From the holdings of the Department of Special Collections , Princeton University Library, Princeton, NJ

These documents are either in the public domain, or can only be used for educational and research purposes (Fair use) as per U.S. Copyright law (text below). By accessing this file, all users agree that their use falls in the public domain or within fair use as defined by the copyrightlaw. We do not charge any permission or usage fees for the publication of images of material inour collections, including those provided by Princeton University Library, via our website, catalog, or directly from staff.For more information regarding copyright, credit, and citation please review the information onthe Copyright, Credit and Citation Guidelines page of our website.
https://library.princeton.edu/special-collections/policies/copyright-credit-and-citation-guidelines. If you have any questions, including how to order higher-quality images from the collection, contact Special Collections Public Services via our Ask Us! Form. https://library.princeton.edu/special-collections/ask-us

---

Doctoral Dissertations, 2011-2026 Philosophy

2014

Modular Theory and Spacetime Structure in QFT

Swanson, Noel Waldron

Princeton, NJ : Princeton University

---

http://arks.princeton.edu/ark:/88435/dsp012v23vw619

U.S. Copyright law: The copyright law of the United States (Title 17, §108, United States Code) governs the making of photocopies or other reproductions of copyrighted material. Under certain conditions specified in the law, libraries and archives are authorized to furnish a photocopy or other reproduction. One of these specified conditions is that the photocopy or reproduction is not to be used for any purpose other than private study, scholarship, or research. If a user makes a request for, or later uses, a photocopy or reproduction for purposes in excess of fair use that user may be liable for copyright infringement. This institution reserves the right to refuse to accept a copying order if, in its judgment, fulfillment of the order would involve a violation of copyright law. https://www.copyright.gov/title17/title17.pdf

Modular Theory and Spacetime Structure in QFT

Noel Waldron Swanson

A Dissertation
Presented to the Faculty
of Princeton University
in Candidacy for the Degree
of Doctor of Philosophy

Recommended for Acceptance
by the Department of
Philosophy
Adviser: Hans Halvorson

September 2014

© Copyright by Noel Waldron Swanson, 2014.

All rights reserved.

Abstract

In quantum field theory (QFT), physical quantities are represented by a net of operator algebras indexed by spacetime region. Each local algebra comes equipped with a state-dependent pair of modular operators, (J, Δ), encoding vital information about the algebra's internal structure, the local state, and the relational structure of the net. These operators and the associated structure theory are the subject of *Tomita-Takesaki modular theory*. Despite its importance, the physical significance of modular theory remains murky, and philosophers of QFT have largely ignored it. This dissertation aims to help close the gap. It explores a series of deep connections between modular theory and spacetime geometry established by the Bisognano-Wichmann theorem (which identifies the modular operators attached to spacelike wedge regions in the vacuum sector with certain generators of the Poincaré group).

Chapter 1 orients my project within the current philosophical literature on the foundations of QFT. I advocate a cosmopolitan stance towards the subject, incorporating tools from Lagrangian, constructive, and algebraic field theory alike. Adopting this stance, chapter 2 introduces the core mathematical ideas behind modular theory and discusses several important physical applications centered around the Bisognano-Wichmann theorem.

Chapter 3 examines how modular theory can be used to give new geometric insight into the parity-charge-time (PCT) theorem. I argue that the key to understanding the mysterious connection between spatiotemporal orientation and charge structure entailed by the theorem lies in recognizing PCT symmetry as a single global reflection of quantum statespace.

Chapter 4 turns to the vexed problem of localization in QFT. Modular theory tells us that at a fundamental level the world is thoroughly entangled. In spite of this, via the split/nuclearity conditions, modular theory also supplies critical tools needed to help explain approximately-localized emergent entities like particles and chairs.

The final chapter contains a critique of the Connes-Rovelli thermal time hypothesis, a proposal to solve the problem of time in quantum gravity using modular theory. Although I conclude that the current proposal cannot provide a coherent gauge-free description of

physical change in generally covariant settings, it does highlight an intriguing connection between modular structure and local dynamics in standard QFT.

Acknowledgements

The present work reflects the philosophical input of a number of individuals who provided helpful comments, criticism, and conversation throughout the writing process: Dave Baker, Jeff Barrett, Thomas Barrett, Gordon Belot, Jeremy Butterfield, Shamik Dasgupta, Ben Feintzeig, Richard Healey, Robbie Hirsch, Michaela McSweeney, Joe Rachiele, Laura Ruestche, Charles Sebens, Dimitri Tsementzis, David Wallace, Jim Weatherall, and Porter Williams. I would also like to thank audiences at Pittsburgh, Irvine, Michigan, and Princeton for feedback on early versions of the central chapters. Above all, I would like to express my deep appreciation and gratitude to my adviser, Hans Halvorson, for his expert guidance over the past five years. This project would not have been possible without him. Finally, I wish to thank my family and friends for their love and support.

To my parents.

Contents

Abstract iii

Acknowledgements v

List of Figures x

1 Cosmopolitan QFT **1**

1.1 Three Variations on QFT 3

1.2 Rival Paradigms? 14

1.2.1 Effective Field Theories 17

1.2.2 Inequivalent Representations 20

1.3 Towards a More Cosmopolitan Worldview 23

Appendix A: Operator Algebras 28

Appendix B: The Haag-Kastler Axioms 34

Appendix C: The Cutoff Reeh-Schlieder Theorem 37

2 Modular Theory for Philosophers **40**

2.1 The Mathematics 41

2.1.1 The KMS Condition 44

2.1.2 Connes's Classification 47

2.1.3 Hilbert Space Geometry 49

2.2 The Physics 51

2.2.1 The Bisognano-Wichmann Theorem 52

2.2.2 Modular Theory and Spacetime Structure 56

3 PCT Symmetry and Statespace Geometry 61

3.1 Spacetime Orientation 65

3.2 Superselection Rules and Antimatter 70

3.3 The PCT Theorem 72

3.4 Dissecting the Theorem 76

3.4.1 The Lie-Jordan Structure and Statespace Reflection 76

3.4.2 Parity and Time Reversal on Wedges 79

3.4.3 Charge Conjugation 84

3.4.4 Summary 88

3.5 Greaves's Account 89

3.6 Classical or Quantum? 95

4 The Type Question 101

4.1 The Type III Property 104

4.2 Taming Entanglement 112

4.2.1 Nuclearity and the Split Condition 113

4.2.2 Local Charge Structure 117

4.3 Spacetime State Realism 125

4.3.1 Cutoffs and the Type III Property 130

4.3.2 Split Localization to the Rescue? 134

5 Can Quantum Thermodynamics Save Time? 137

5.1 The Thermal Time Hypothesis 139

5.2 Thermal Time and Proper Time 145

5.3 The Classical Limit 152

5.4 Conceptual Challenges 156

5.5 Modular Dynamics 161

Bibliography **164**

List of Figures

1 The analogy between $\mathfrak{M}$ and the complex plane 49

2 Relative location of real subspaces generated by ω, ϕ. 51

3 A spacelike wedge 53

4 The P_1 reflection implemented by AdJ 81

5 A split inclusion $\Lambda = \{\mathfrak{R}(\mathcal{D}), \mathfrak{R}(\tilde{\mathcal{D}}), \Omega\}$ 115

Chapter 1

Cosmopolitan QFT

Quantum field theory (QFT) is a branch of physics which unifies core ideas from quantum mechanics and special relativity.[1] It provides the framework for the Standard Model of particle physics, a theory which gives a near complete description of all known fundamental matter and forces in our world apart from gravity. The Standard Model has achieved extraordinary empirical success; it is probably the most thoroughly tested scientific theory in human history.[2] Nonetheless, at present our mathematical and conceptual understanding of QFT is still in its infancy.

Not only does QFT inherit many of the same interpretational problems that plague ordinary quantum mechanics, but even before we can grapple with issues like the measurement

[1] Throughout, we will focus on QFTs defined over flat Minkowski spacetime. We can represent Minkowski spacetime using a smooth 4-dimensional manifold, $\mathcal{M}$, equipped with a flat Lorentzian metric, η, of signature $(+,-,-,-)$. The metric determines the spatiotemporal distance between points and defines the causal lightcone structure of Minkowski spacetime. At every point, p, there is a double-lobed cone, whose boundary is formed by the possible trajectories of massless photons passing through p. A vector at p is said to be *lightlike* if it lies along the edge of the cone, *timelike* if it lies in the interior, and *spacelike* if it lies outside the cone. (More formally, the metric determines a unique isomorphism between the tangent bundle $T\mathcal{M}$ and its dual space $T^*\mathcal{M}$, assigning every 4-vector a unique covector, $x^a \mapsto x_a$. A 4-vector is lightlike iff $x^a x_a = 0$, timelike iff $x^a x_a > 0$, and spacelike iff $x^a x_a < 0$.) On the usual interpretation, the laws of relativistic mechanics require matter fields to propagate inside lightcones. The *Poincaré* group, $\mathcal{P}$, represents the group of continuous metric-preserving symmetries of Minkowski spacetime. Intuitively, it includes translations, rotations, Lorentz boosts, and reflections. The connected component ($\mathcal{P}$ minus the reflections) is denoted $\mathcal{P}^{\uparrow}_{+}$.

[2] An oft-cited example: the anomalous magnetic moment of the electron has been shown to agree with the Standard Model predictions to within one part in a billion (Hanneke et al. (2011)).

problem, we need to figure out how to write down mathematically rigorous versions of the theory in the first place. The textbook Lagrangian approach to QFT familiar to most physicists is a grab bag of different mathematical ideas. While astonishingly fruitful, it has a number of troubling holes and formal inconsistencies which require ad hoc patches. Axiomatic approaches like algebraic and constructive field theory offer greater rigor, but have yet to reproduce the successes of the Standard Model. In many ways, the current state of play is akin the the "wild west" days of early quantum mechanics back in the 1920s before the publication of von Neumann's *Grundlagen* established a unified mathematical language for non-relativistic quantum theory.

It is the relativistic analogue of von Neumann's project that a small but devoted group of physicists, mathematicians, and philosophers have spent the last fifty years working on. Indeed the task has proven to be far more challenging than anyone initially imagined. It should be emphasized, however, that the project is not just one of foundational, academic interest. With the recent discovery of the Higgs boson, particle physics is at an interesting developmental stage. On one hand, the discovery cements the last piece of the Standard Model firmly in place. On the other, it now seems increasingly unlikely that major clues about physics beyond the Standard Model will emerge from accelerator experiments anytime soon. With the eventual goal of developing a working theory of quantum gravity down the line, it is incredibly important at this juncture for us to reexamine the foundational underpinnings of QFT.

The present chapter serves to orient my dissertation within this larger foundational project. In §1.1, I give a brief overview of three of the main approaches to the foundations of QFT — Lagrangian QFT, constructive QFT, and algebraic QFT. Recently, a number of philosophers have characterized these approaches as inequivalent, rival research programs, arguing that it is important to determine which approach is best before the task of interpreting QFT can begin. I strongly disagree. In §1.2, I examine this philosophical debate focusing on a recent exchange between David Wallace (who champions Lagrangian QFT)

and Doreen Fraser (who defends algebraic QFT). While each approach emphasizes different aspects of QFT, I contend that there is in fact no deep theoretical tension between them. Ultimately we hope to see convergence between these programs, although at present there are still significant technical and conceptual hurdles to overcome. Given the limitations of our current knowledge and the need for creative new ideas, I argue in §1.3 that philosophers should adopt a *cosmopolitan* approach to the foundations of QFT which aligns itself more closely with scientific practice. I conclude by discussing how the main subject of my dissertation, Tomita-Takesaki modular theory, fits within this cosmopolitan picture. In particular, I think that modular theory stands in a unique position to help bridge the divide between different foundational approaches. I explain why, highlighting some of the central topics which will be taken up in detail in subsequent chapters.[3]

1.1 Three Variations on QFT

The proliferation of various mathematical formalisms for QFT has been a source of significant confusion in the philosophy of physics literature. Unlike the situation in other corners of physics like classical mechanics, statistical mechanics, special and general relativity, and non-relativistic quantum mechanics, we do not have a canonical mathematical framework for QFT that is both internally consistent and empirically adequate. Instead, we are faced with an array of different mathematical tools and techniques of varying degrees of rigor which frequently clash with each other. In this section, we'll survey three of the most general frameworks: Lagrangian QFT, constructive (Wightman) QFT, and algebraic QFT.

[3]Material from this chapter was originally presented at colloqiua at the University of Minnesota and Carnegie Mellon University in January 2014 and the University of Delaware in February 2014.

Lagrangian QFT

Pragmatically speaking, by far the most successful approach to QFT has been Lagrangian field theory.[4] This is the mainstream approach taught to every physics graduate student and the one used to frame the Standard Model in its present incarnation. The guiding idea behind Lagrangian QFT is to take tools from classical field theory and adapt them to the quantum realm. We typically begin with a classical field system whose dynamically possible configurations are encoded in a single master function called the Lagrangian. We then apply a series of quantization rules to translate this system into a quantum field theory. An assortment of approximation techniques collectively known as perturbation theory can then be used to extract information from the quantum model. Though far from ideal foundationally-speaking, these Lagrangian methods have proven to be tremendously powerful and have found many fruitful applications, both in particle physics as well as other domains.

Broadly speaking, a *field* is an extended object with infinitely many degrees of freedom which are characterized by an assignment of mathematical values to each point in spacetime. These values can be any sort of mathematical object you like, real numbers, complex numbers, vectors, tensors, spinors, etc.[5] In the case of quantum fields, the field values are mathematical objects called *operators*. These operators act as linear transformations on a complex vector space called a *Hilbert space*.[6] As in non-relativistic quantum mechanics, we employ Hilbert spaces to represent the space of possible states the quantum system. In the relativistic case, states are usually taken as temporally extended rather than instantaneous. We can thus view Hilbert space as a geometric model of the space of physically possible

[4] See Weinberg (2005), Peskin and Schroeder (1995) for standard treatments.

[5] Fields are often modeled geometrically as sections of a fiber bundle over spacetime. For present purposes we can rely on a less sophisticated view of fields as functions $\phi : \mathcal{M} \to V$, where $\mathcal{M}$ represents spacetime and V is an appropriate value space (e.g. $\mathbb{R}$, $\mathbb{C}$, or some other vector space).

[6] Hilbert spaces in QFT are assumed to be *separable*, meaning that they admit a countable orthonormal basis.

worlds (or world histories) described by the theory. So the "field" part of QFT refers to a field of transformations, indexed by spacetime coordinates, acting on this modal space.

The physical significance of this field of operators is somewhat mysterious. While mathematically similar to the operators used in standard quantum mechanics to represent observable physical quantities like energy, spin, charge, etc., the field operators in QFT do not represent such quantities directly. Instead, we construct observables out of algebraic combinations of the field operators. If there is a gauge symmetry present (as is the case in all theories like the Standard Model), only those algebraic combinations of the fields which are invariant under the gauge action are deemed to have direct physical significance. To put it metaphorically, the field operators are superfluous descriptive fluff, much like the vector potential in classical electromagnetism.[7] In addition, while most observables are represented by bounded operators, the field operators are unbounded, making their technical manipulation quite tricky.[8] For now, we can ignore these interpretive worries and focus on the mathematical viability of the field concept.

Since Lagrangian field theory is really a cluster of different computational techniques, it does not give us a clean set of mathematical conditions on what counts as a model of QFT. This makes it challenging to prove rigorous general theorems. Studying field theories in the Lagrangian approach usually proceeds on a model by model basis. In addition, there are internal consistency problems that arise at the level of individual model building. These can be broken into three broad classes.

First there are problems with the quantization procedures. Standard canonical quantization techniques treat non-linear interaction terms as small perturbations of a free field system. For non-abelian gauge theories like those used to model the weak and strong nuclear force, this assumption fails. There are more powerful techniques like path integral

[7]There is still ongoing debate about whether or not this is the correct interpretation of gauge symmetries in QFT, but it's the consensus view amongst a healthy majority of working physicists and philosophers alike, and we'll take it on board for now.

[8]An unbounded operator $T : \mathcal{H} \to \mathcal{H}$ is a linear map defined on a dense linear subspace $D \subseteq \mathcal{H}$ (typically D is not the full Hilbert space). T is bounded from below if $\langle Tx|x\rangle \geq -c||x||^2$ for some real constant c. T is bounded if both T and $-T$ are bounded.

quantization, but these require choosing a measure over an infinite dimensional path space, a procedure which is only mathematically well-defined in certain special cases. In practice, these complications are typically ignored.[9] These technical issues aside, once we settle on a particular quantization method, there is still the vexing problem of inequivalent representations. Since field systems have infinitely many degrees of freedom, quantization procedures typically do not yield a unique output. Instead, given a single classical input, we get an infinite family of physically inequivalent quantum models. How to make sense of this explosion of possibilities has been a topic of much debate.[10] We'll look at this problem in more detail in §1.2.2.

A second set of issues concern whether or not the output quantum models are mathematically consistent. The process of extracting physical predictions from a quantized model requires computing functional integrals which include contributions from physical processes at both extremely short distances/high energies and extremely long distances/low energies. Both can cause the integrals to blow up, yielding senseless infinities. It turns out that this is a generic problem in interacting QFTs. Even when the practical challenges posed by ultraviolet (UV) and infrared (IR) divergences can be tamed using cutoffs and renormalization group techniques, a number of conceptual problems they generate continue to linger. For example, the standard picture of particle interactions models collisions using a unitary S-matrix which maps free incoming states to free outgoing states. Haag's Theorem shows that such a unitary operator does not exist unless the field is free at all times.[11] Ideally one would like to explain the success of the standard scattering picture by replacing the assumption of

[9]More mathematically rigorous quantization procedures exist such as geometric quantization and deformation quantization, although these take their cue from the mathematics of constructive and algebraic QFT, respectively, and are very difficult to implement in concrete models. A more thorough understanding of various quantization techniques, their domains of applicability, and how they are interrelated is a major open problem in the foundations of quantum theory. See Landsman (1998) for a thorough treatment of more rigorous quantization methods.

[10]Ruetsche (2011a) examines both the interpretive challenges and numerous explanatory applications of inequivalent representations in infinite quantum systems.

[11]See Haag (1996, Ch. III.3). Also Earman and Fraser (2006) for a philosophical account of the foundational implications of Haag's theorem.

free incoming and outgoing states with asymptotically free states, but at this stage we only know how to implement this idea rigorously in certain restricted situations.[12]

Finally, there are questions about whether or not the various perturbative techniques we employ yield an accurate approximation of the underlying physics. Dyson gave a famous heuristic argument showing that the perturbation series in quantum electrodynamics (QED) does not converge. There is an ongoing debate about various formalized versions of this argument.[13] If Dyson's argument does go through, it might still be the case that the perturbation series is asymptotic to a well-defined limiting QFT. Indeed we have excellent empirical reasons to trust perturbative QED. It clearly gives us an approximation of some rich underlying physics. The problem is that without a better theoretical understanding of the non-perturbative situation we can't say very much about what this underlying physics looks like. Recent triviality results for an important interacting model called ϕ^4-theory highlight the problem. Independently, Aizenman and Frölich have proven that in $d > 4$ spacetime dimensions the limit theory is actually a free field theory despite the fact that the perturbation series appears to describe non-trivial interactions. There are similar results for ϕ^4-theory in $d = 4$ although these rely on certain technical assumptions which need further scrutiny.[14] (In $d < 4$ the the theory is non-trivial.) These results indicate that we still do not fully understand when we can trust the perturbative techniques employed by Lagrangian QFT.

In order to provide a mathematically precise answer the question "what is a model of QFT" and to address the internal consistency problems discussed above, various axiomatic approaches have been developed. The two most well established frameworks are constructive field theory and algebraic field theory. Both start from standard Lagrangian models and attempt to frame precise, general principles that any physically reasonable QFT must obey.

[12]The current best approach is Haag-Ruelle scattering theory (Haag, 1996), but these techniques only work if there are both upper and lower bounds on the particle masses. While the former assumption (i.e. a mass gap) is a physically reasonable requirement, the latter is not.

[13]Summers (2012a), §3.3.

[14]See Fernandez and Sokal (1992) for an exhaustive treatment.

They focus on slightly different parts of the theory, however, and as a result, they differ in their respective strengths and weaknesses.

Constructive (Wightman) QFT

The axioms for constructive field theory were first proposed by Wightman and Gårding in the 1950s.[15] They hew closely to Lagrangian field theory, describing a system of covariant operator-valued quantum fields acting on a single Hilbert space. A series of no-go results show that due to UV divergences, field operators cannot be assigned to individual spacetime points. We can only coherently talk about the field configuration in extended regions of spacetime.[16] Nonetheless, the axioms ensure that these regions can always be made arbitrarily small and compact. Formally this requires representing the fields as tempered operator-valued distributions, $\phi(f)$, where f is an arbitrary smooth test function of compact support on the background spacetime manifold $\mathcal{M}$. As in Lagrangian QFT, the field operators are generally unbounded, and act on fixed Hilbert space whose unit rays represent physically possible (global) states of the system.

A model of the Wightman axioms consists of a 4-tuple $(\{\phi\}, \mathcal{H}, U, \Omega)$, where $\{\phi\}$ is a family of fields acting on a Hilbert space $\mathcal{H}$ as described above, U is a representation of Poincaré group (or more generally the translation subgroup) in terms of unitary operators acting on $\mathcal{H}$,[17] and Ω is a unique translation-invariant vector in $\mathcal{H}$ representing the vacuum state. There are three central axioms:

Covariance The fields $\{\phi\}$ must transform covariantly under Poincaré transformations.

If $U(a, \Lambda)$ is a unitary operator representing a combined translation, a, and Lorentz

[15] Streater and Wightman (1989) and Jost (1965) are classic sources. Throughout the remaining chapters we will use constructive QFT as synonymous with Wightman QFT. There is actually a host of different constructive approaches which are all linked by the common starting point of unbounded, gauge-dependent field operators. See Summers (2012a) for an excellent survey of both historical topics and current developments within this expansive field.

[16] See Halvorson and Müger (2006), §6.1 for various no-go results concerning field operators at points.

[17] Unitary operators are linear isometries of $\mathcal{H}$ which preserve all transition probabilities. Wigner's theorem asserts that any symmetry of a quantum system must be implemented by unitary (or antiunitary) operators. More on this later.

boost, Λ, then

$$U(a,\Lambda)\phi(f(x))U(a,\Lambda)^{-1} = \phi(f(\Lambda x + a)).$$

This ensures that the theory is invariant under the symmetries of Minkowski spacetime.

Microcausality Fields must either commute or anti-commute at spacelike separation (i.e. either $\phi(f)\phi(g) - \phi(g)\phi(f) = 0$ or $\phi(f)\phi(g) + \phi(g)\phi(f) = 0$ when f an g have spacelike separated support). This condition enforces no-signaling constraints and is closely tied to the spin-statistics theorem. In order for a model to satisfy the remaining axioms, fields with integer spin must commute and fields with half integer spin must anti-commute. The former condition gives rise to symmetric Bose-Einstein statistics while the latter give rise to antisymmetric Fermi-Dirac statistics.

Spectrum Condition The spectrum of the energy-momentum operator P^μ must be contained in the closed forward lightcone in momentum space. This ensures that the energy is positive all Lorentz frames and that the vacuum state is a ground state. It is a crucial assumption for securing the stability of matter.

In addition there are several technical assumptions concerning the domain of definition of the field operators and the cyclicity of the vacuum state.[18]

The first fully rigorous proofs of a number of central theorems like the PCT and spin-statistics theorems were achieved in the setting of constructive QFT. Additionally, constructive techniques give us a powerful set of tools for building non-perturbative models of QFT. One such tool is the Wightman reconstruction theorem which allows us to recover a full model from its set of n-point vacuum correlation functions $\langle\Omega|\phi(f_1)\ldots\phi(f_n)\Omega\rangle$. Despite these strengths, the formalism does have certain weaknesses. First, the axioms concern gauge-dependent field operators. Since these do not directly represent physical quantities, it makes the physical interpretation and justification of the axioms difficult. Second, many of the proofs in constructive field theory rely on the technique of analytically extending func-

[18]See Streater and Wightman (1989, Ch. 3.1) for a full list of axioms.

tions of the field operators, a procedure which lacks clear physical or explanatory content. Third, a series of no-go results due to Strocchi indicate that the axioms cannot accommodate local gauge theories like QED which employ field operators that cannot be localized in arbitrarily small, compact regions.[19]

Algebraic QFT

Rather than assumptions about field operators, algebraic QFT (AQFT) starts with axioms for gauge invariant observables.[20] On this picture, each region of spacetime is associated with a collection of physical quantities called the local algebra of observables. These correspond to quantities which can be measured by experiments causally confined to that region. At the outset, we do not assume that these quantities are represented by operators acting on a particular Hilbert space. Instead, each collection comes equipped with an intrinsic algebraic structure (usually that of a C^* or von Neumann algebra) which can be used to directly describe functional relations between various quantities.[21]

The *Haag-Kaslter axioms* ensure that the assignment of algebras to regions satisfies consistency conditions imposed by relativistic causality and quantum mechanics. Letting $\mathfrak{A}(\mathcal{O})$ denote the local algebra for region $\mathcal{O}$, the first axiom ensures that any observable measurable in $\mathcal{O}$ is measurable in any larger region containing $\mathcal{O}$:

Isotony If $\mathcal{O}_1 \subseteq O_2$, then $\mathfrak{A}(\mathcal{O}_2)$ contains $\mathfrak{A}(\mathcal{O}_1)$ as a subalgebra.

The isotony axiom endows the collection of local algebras with the structure of a net. Taking the upwards inductive limit of this net generates the *quasi-local* algebra of observables,

[19] See Strocchi (1993) for full details. Horuzhy (1990) contains a succinct overview in the appendix.

[20] The locus classicus is Haag (1996). An excellent recent philosophical survey of AQFT focusing on the problem of inequivalent representations is Ruetsche (2011a). For a compact, technical introduction see Brunetti and Fredenhagen (2004). Halvorson and Müger (2006) is a thorough mathematical treatment of the subject from a philosophical viewpoint.

[21] Such algebras mimic the algebraic properties of bounded linear operators acting on a Hilbert space. See Appendix A for more details.

usually denoted by $\mathfrak{A}$. The quasi-local algebra contains all local observables along with all observables that can be approximated by sequences of local observables.[22]

There are also algebraic analogues of covariance and microcausality:

Covariance For each Poincaré transformation (a, Λ), there exists an automorphism, $\alpha_{(a,\Lambda)}$, of the quasilocal algebra, $\mathfrak{A}$, that acts covariantly on localized subalgebras $\mathfrak{A}(\mathcal{O}) \subset \mathfrak{A}$:

$$\alpha_{(a,\Lambda)}(\mathfrak{A}(\mathcal{O})) = \mathfrak{A}(\Lambda\mathcal{O} + a).$$

Microcausality If $\mathcal{O}_1$ and $\mathcal{O}_2$ are spacelike separated regions, then every observable in $\mathfrak{A}(\mathcal{O}_1)$ commutes with every observable in $\mathfrak{A}(\mathcal{O}_2)$.

Algebraic covariance is a weaker, more natural assumption than Wightman covariance since it it possible to build a covariant net of observables out of non-covariant Wightman fields. Assuming that the field operators are gauge quantities, such models should intuitively count as physically reasonable relativistic theories. Algebraic microcausality has a similar advantage over its Wightman counterpart. In $d < 4$ spacetime dimensions, it allows for interesting alternatives to standard statistics such as braid group statistics and parastatistics.

States are represented abstractly as positive, normalized linear functionals, $\omega : \mathfrak{A} \to \mathbb{C}$. The value of the functional $\omega(A)$ is interpreted as the expectation value for the observable $A \in \mathfrak{A}$ in state ω. Given a state, the *Gelfand-Naimark-Segal (GNS) construction* determines a unique Hilbert space representation of the observable algebra.[23] A representation is simply a mapping from the algebra into bounded operators on some Hilbert space, $\pi : \mathfrak{A} \to \mathcal{B}(H)$, which preserves all of the relevant algebraic structure encoded in $\mathfrak{A}$.[24] Many of these representations carry direct physical significance. Working in a concrete representation gives us additional topological tools which can be used to define important physical quantities

[22]The relevant topology here is the uniform topology induced by the canonical C^*-algebra norm. See Appendix A for more details.

[23]See Kadison and Ringrose (1997a, Thm. 4.5.2)

[24]In particular π must be a $*-$homomorphism.

including temperature, stress-energy, superselected gauge charges, order parameters for systems near phase transitions, and in Fock space theories, the global particle number operator. Unlike constructive or Lagrangian field theory, AQFT typically makes use of multiple unitarily inequivalent Hilbert space representations. It is this flexibility that has allowed the approach to provide deep insight into a range of phenomena including quantum phase transitions, spontaneous symmetry breaking, charge superselection rules, particle statistics, and the concept of antimatter.

The GNS representation associated with the vacuum state ω_0 plays a particularly central role. In the vacuum representation, spatiotemporal translations can be represented by unitary operators. The infinitesimal generator of the translation subgroup corresponds to the energy-momentum observable, P^μ, and just as in the Wightman picture, we require the spectrum condition to hold.

The main axioms are frequently supplemented with additional constraints. Many of them are designed to select a privileged set of physically relevant representations. For example, DHR superselection theory, which describes theories with strongly localized charges like quantum chromodynamics (QCD), begins with the assumption that only representations that differ from the vacuum representation within some compact region are physically permissible. Another technical assumption which will be important in the next section is the following:

Weak Additivity The quasi-local algebra $\mathfrak{A}$ can be generated by taking spatiotemporal translations of any local algebra, $\mathfrak{A}(\mathcal{O})$.[25]

As a consequence of weak additivity, any physical observable can be constructed out of observables associated with arbitrarily small compact regions. This axiom is usually motivated on the grounds that there shouldn't be a smallest lengthscale in the theory. (As we'll go on to see, however, it's not clear that this should always be the case.)

[25]Strictly speaking this requires that the local algebras are von Neumann algebras. If we start with a C^*-algebra, $\mathfrak{A}(\mathcal{O})$, we can define a local von Neumann algebra, $\pi(\mathfrak{A}(\mathcal{O}))^-$, by taking the closure in the weak operator topology on the GNS representation associated with the global state. By the von Neumann bicommutant theorem $\pi(\mathfrak{A}(\mathcal{O}))^- = \pi(\mathfrak{A}(\mathcal{O}))''$.

Bringing all of these pieces together, we can think of a model of AQFT as a 4-tuple $(\{\mathfrak{A}(\mathcal{O})\}, \{(\pi, \mathcal{H})\}, \omega_0, U)$, where $\{\mathfrak{A}(\mathcal{O})\}$ is a collection of algebras associated with spacetime regions $\mathcal{O}$, $\{(\pi, \mathcal{H})\}$ is collection of physically relevant Hilbert space representations of the quasi-local algebra, ω_0 is the vacuum state, and U is a unitary representation of the Poincaré group (or more generally the translation subgroup) acting on the vacuum representation $(\pi_0, \mathcal{H}_0)$.

Because the algebraic approach has more physically transparent initial assumptions, many general theorems in QFT receive their sharpest, most perspicuous formulation within the setting of AQFT. It largely for this reason that AQFT has received a fair amount of attention from philosophers of physics as of late. At the same time, it turns out to be quite difficult to build models of the Haag-Kastler axioms directly. Usually the strategy is to build a constructive model and show that it also satisfies the Haag-Kastler axioms.

To date, the following classes of QFTs are under full mathematical control using a combination of constructive and algebraic methods:

- Free theories (all dimensions)
- Super-Renormalizable theories (d=2,3)
- Conformal theories (d=2)
- Topological theories (d=3)
- Integrable heories (d=2)

Notably absent from this list are Yang-Mills type gauge theories in $d = 4$ which comprise the Standard Model. By far the biggest challenge for both constructive and algebraic QFT is recovering the predictive and explanatory success of the Standard Model. As we have briefly noted, there are reasons to think that the Wightman framework needs to be modified in order to incorporate local gauge theories. Since the Haag-Kastler axioms do not make any direct assumptions about the structure of gauge fields, there is hope that the framework is

already general enough to capture such theories. Alas, at the moment all we have are some enticing theoretical clues and a promissory note.

1.2 Rival Paradigms?

With three different versions of QFT on the table, a natural foundational question arises: how are these three approaches related to each other? A number of philosophers have weighed in on this question. Recently, there has been a notable exchange between David Wallace and Doreen Fraser on the subject. Both agree that these different approaches (or at least the Lagrangian and axiomatic approaches) represent rival theoretical alternatives. But that's where the agreement stops. Here's Fraser:

> QFT presents a genuine example of the underdetermination of theory by empirical evidence. There are variants of QFT which are empirically indistinguishable yet support different interpretations.[26]

Fraser goes on to argue that before the philosophical work of interpreting QFT can go forward, we need to decide which formulation should be subject to interpretation. Fraser champions AQFT. Wallace on the other hand, contends that philosophers should engage directly with Lagrangian QFT. Here's Wallace:

> It is no longer appropriate — if it ever was — to see axiomatic QFT as the proposed mathematically rigorous version of conventional QFT. The two are better understood as *rival research programs*, trying in different ways to resolve the problem of renormalization.[27]

Before we wade into the fray, it will be useful to have a guiding picture. Since algebraic and constructive field theory have precise axioms, we can ask well-posed mathematical questions about their relationship. Although they have a large intersection of common models,

[26] Fraser (2009)
[27] Wallace (2011)

including all those listed in the last section, the two approaches are not equivalent (at least in their standard form). Given a constructive model $(\{\phi\}, H, U, \Omega)$ we can always construct a net of associated observable algebras by taking arbitrary polynomials of the field operators. The local polynomial algebra, $\mathfrak{P}(\mathcal{O})$, consists of all polynomials of the field operators whose test functions have support in $\mathcal{O}$. The resulting net satisfies all of the Haag-Kastler axioms except the requirement that the local algebras be C^*-algebras. There are a number of sufficient conditions known to ensure that this additional requirement is met. The most physically well-motivated place bounds on the local field strength in states with finite energy content.[28] Such bounds are satisfied in all known models of physical interest. Without them it is possible to construct pathological models where the local field strength can be arbitrarily high while the energy is arbitrarily low.

Thus it appears that with some mild physicality assumptions in place, constructive QFT becomes a subtheory of AQFT. In the opposite direction, much less is known. Since the Haag-Kastler axioms don't make any assumptions about underlying field operators, we can construct models that patently violate the Wightman axioms by positing field systems with unusual properties which still give rise to well-behaved nets of observable algebras. Since these axioms describe gauge-dependent entities, however, determining the physical content of the resulting models can be tricky. For example it is possible to build algebraic field theories with mixed-symmetry parastatistics by positing field operators that obey non-standard trilinear commutation relations at spacelike separation. On the face of it, such models directly violate the Wightman axioms. As it turns out, though, a general equivalence theorem shows that it is always possible to construct a field system with ordinary bilinear commutation relations and an additional global non-abelian gauge symmetry which has satisfies the Wightman axioms and has the same gauge-invariant physical content as the original parastatistical theory.[29] There are, however, other physically interesting models of AQFT, such

[28] These are known as H-bound assumptions. See Borchers and Yngvason (1992). Necessary conditions are still elusive.

[29] Baker et al. (2014)

as low dimensional theories containing anyons, which do not suffer the same deflationary fate. Along these lines, many physicists hope that the current algebraic approach is broad enough to capture Yang-Mills theories which we know lie outside the current constructive framework due to the Strocchi theorems.

The Lagrangian approach does not have well-defined boundaries. It remains an open possibility that certain extant Lagrangian models do not fit into either axiomatic framework. Checking is next to impossible, however, since Lagrangian methods do not give us a very good picture of the non-perturbative structure of QFT. On the other hand, recent algebraic construction methods have produced interacting QFTs whose dynamics cannot be specified by a Lagrangian function on the space of field configurations.[30] Thus the algebraic approach is broader in some sense. (It remains to be seen, however, to what extant these models are physically well-behaved.)

There are a number of open questions here. First, we do not have a clear idea about where the exact boundary of Lagrangian QFT is. Second, even though the boundaries of the axiomatic approaches can be drawn more cleanly, we lack a full understand what necessary and sufficient conditions need to be in place in order to easily translate back and forth between the algebraic and constructive QFT. Third, and most importantly, while mathematically precise, the current axioms are only provisional. We are still in the process of tinkering with each framework, looking for possible extensions, supplemental assumptions, and other modifications. For example, we know that both algebraic and constructive QFT need to be modified in order to make sense of QFT in curved spacetimes. Given the current state of flux, in order to make an argument for inequivalence, Fraser and Wallace need to point to a "essential" assumptions of each approach which directly conflict with one another.

[30] Summers (2012a), §6-7.

1.2.1 Effective Field Theories

Wallace thinks he has one. Specifically, he claims that axiomatic QFT is fundamentally incompatible with the concept of an *effective field theory*, an arguably central component of Lagrangian QFT.

Given a fundamental QFT which describes physics at all lengthscales, suppose we want to explore physics in the long-distance limit. Renormalization theory give us a set of sophisticated coarse graining tools for extracting stable limit regularities from the fundamental theory. In the best cases, it turns out that the long-distance physics is largely insensitive to the specific details of the underlying theory. What is more, renormalization gives us a kind of explanation for when and why this occurs. This is a concrete instance of multiple realizability. In fact renormalization techniques have been applied to a wide range of similar problems in other domains including chemistry, systems biology, and economics.

The effective field theory picture, turns this scenario on its head. Say we don't yet know what the fundamental theory is. It might not even be a QFT at all, but rather a version of string theory, quantum loop gravity, or something else entirely. We still have good empirical and theoretical reasons to suspect that the long-distance limit theory is a QFT. We can proceed by guessing an effective field theory — writing down a limit-QFT with an arbitrary, floating cutoff parameter built-in. We stipulate that the limit-QFT only accurately describes phenomena down to some finite lengthscale, λ. We then make a further assumption that the theory only includes interaction terms which are renormalizable, that is, insensitive to the details of the cutoff and the underlying short-distance physics. We can then test the theory's empirical predictions and hopefully confirm it. On this view, an effective field theory is an intrinsically approximate beast, one that we can be confident in even if we don't understand what lies underneath it.

Most physicists think that the Standard Model is an effective field theory of this kind. Over a series of papers, Wallace develops this point into the following argument:

P.1 The Standard Model is best viewed as an effective field theory.

P.2 Axiomatic QFT is fundamentally incompatible with an effective field theory picture.

P.3 The Lagrangian formulation of the Standard Model has achieved remarkable empirical and explanatory success.

P.4 Wilsonian renormalization has put Lagrangian field theory on secure mathematical footing.

$\therefore$ Philosophers of QFT should focus on the Lagrangian formulation of the Standard Model.

In support of the crucial premise P.2, Wallace points to the weak additivity axiom in AQFT. Recall that the axiom stipulates that the quasi-local algebra, $\mathfrak{A}$, can be generated by translations of any local algebra, $\mathfrak{A}(\mathcal{O})$. The important part is the assumption that $\mathcal{O}$ can be any region, even arbitrarily small, compact ones. So it looks like AQFT is in principle incompatible with a short-distance cutoff since weak additivity requires that the local algebras of observables are non-trivial at all scales.

Looks can be deceiving. Once we get clear on how weak additivity is actually deployed in proving theorems, we find that the vast majority of deep foundational results obtained by AQFT continue to hold if we replace weak additivity by a suitable cutoff variant. As it turns out, weak additivity is almost exclusively used to prove versions of the Reeh-Schlieder theorem, a central result which tells us that the vacuum state is superentangled across all local algebras. Normally the theorem holds at all length-scales. If instead we assume that weak additivity only holds for regions larger than some cutoff parameter, we can still prove a version of Reeh-Schlieder for the associated local algebras:

Proposition 1.1 (Cutoff Reeh-Schlieder)**.** *Assume that weak additivity holds for algebras associated with regions larger than $\mathcal{O}_\lambda$ for some finite scale parameter λ. If the spectrum condition and microcausality hold, the Reeh-Schlieder property obtains for algebras associated with regions larger than $\mathcal{O}_\lambda$.*[31]

[31] For a precise formulation and proof, see Appendix C.

As we'll see in the next chapter, the Reeh-Schlieder theorem is essential for getting the entire program of modular theory off the ground. Fortunately, the cutoff variant of the theorem is sufficient to establish the geometric properties of the modular operators above the cutoff scale. In addition, Buchholz and Verch have developed new tools called scaling algebras for modeling renormalization group flows within AQFT.[32] While more work in this direction is definitely needed, these techniques have already begun to shed light on the non-perturbative nature of quark confinement and asymptotic freedom in QCD. Renormalization has such a central explanatory role in QFT, that being able to make sense of it should be taken as a constraint on any adequate axiom system.

There are two results which might need to be revised in algebraic models of effective field theories. The first is a classification theorem for local algebras which tells us that in virtually all physically realistic models, the local algebras are a special kind of C^*-algebra called type III_1 von Neumann algebras (see Ch. 2.1.2, 4.1). The main proof of this result uses the fact that in the short distance scaling limit, physically reasonable theories will be asymptotic to conformal field theories.[33] Whether or not this is true will depend on the details of the cutoff. In Ch. 4.3.1, however, we give an argument supporting the conclusion that even in effective field theories, we can expect the local algebras to be at least type III von Neumann algebras.

The second result is the Doplicher-Roberts reconstruction theorem which allows us to recover the structure of an underlying system of field operators and gauge group from assumptions about the charge structure of a given algebraic model.[34] Here it is important to emphasize that it is only the field reconstruction theorem, not the more general DHR analysis of superselection rules which appears to require weak additivity for arbitrary compact regions. (In fact the DHR analysis concerns the details of the physics at the long-distance not short-distance limit.) Thus the insights that DHR provides into charge structure, antimatter,

[32]Buchholz and Verch (1995) We'll discuss the role of scaling algebras in the proof of the type III property in Ch. 4.1.

[33]Fredenhagen (1985a)

[34]Doplicher and Roberts (1990)

and particle statistics still stand. That the Doplicher-Roberts reconstruction theorem would require weak additivity is not entirely surprising since the Wightman framework posits well defined field operators $\phi(f)$ for arbitrary smooth test functions of compact support. This assumption presents a similar undermining problem for the effective field theory picture in the constructive approach.

Fortunately, this problem appears to be of a more practical rather than foundational nature. It is true that many of the primary model-building tools supplied by constructive field theory such as the Wightman Reconstruction Theorem, rely on analytic extension techniques which presuppose well-defined field operators at all length scales. Despite this, such techniques are usually viewed as a pragmatic means of constructing a gauge-invariant net of observables. We know from the algebraic approach a that such nets are compatible with an effective field theory viewpoint, so even if we have to modify our construction methods the problem does not concern the physical content of the theory.

So contra Wallace, AQFT appears to already have the tools needed to make sense of effective field theories. Constructive QFT may require new tools, but nothing on the table so far suggests that there is any deep incompatibility.

1.2.2 Inequivalent Representations

Fraser responds to Wallace in a different way, leveraging her objection into an argument that philosophers should continue to focus primarily on AQFT. She maintains that P.3 and P.4 are overstated. In particular, she argues that there are certain central explanatory patterns in QFT which only the algebraic approach can make sense of. On her view explanatory utility plus mathematical precision along with the provisional nature of various axiomatizations trump Wallace's considerations.

At the heart of her argument is the collection of representations, $\{(\pi, \mathcal{H})\}$, of the quasi-local algebra, $\mathfrak{A}$. As we saw in §1.1, these representations are needed to make sense of certain globally defined quantities which cannot be constructed from the local observables alone.

Because quantum fields have an infinite number of degrees of freedom, these representations are unitarily inequivalent.[35] Thus AQFT employs the concept of inequivalent representations to model systems which differ from one another with respect to global quantities which cannot be measured in any local region. For instance, in the DHR analysis of superselection rules, global states which differ from another with respect to the global charge, Q, are represented by state vectors in different representations. Similar techniques are crucial in the algebraic treatment of spontaneous symmetry breaking, phase transitions, particle statistics, and antimatter.

Fraser argues that since Lagrangian field theory necessarily employs cutoffs, the approach is incompatible with explanations which make use of inequivalent representations. The cutoffs render all representations finite-dimensional, and all such representations are unitarily equivalent according the Stone-von Neumann theorem.

There are three things to say in response. The first is a general point about explanation and idealization. As we know from cases in classical statistical mechanics, it is possible for a finite theory to appeal to an idealized infinite limit for explanatory purposes even when the finite mathematical model is strictly speaking incompatible with the idealization. One way to make sense of this situation is to appeal to a unificationist account of explanation on which the infinite idealization allows for the unification of a diverse array of finite systems. So prima facie, it is not obvious that the mathematical incompatibility that Fraser points to is indicative of an explanatory incompatibility.

Second, and more critically, Fraser really only has half of the story right. Recall that Lagrangian QFTs are afflicted by both ultraviolet and infrared divergences. For the purposes of mathematical consistency, the effective field theory picture only really requires a short distance, ultraviolet cutoff. Infrared divergences can sometimes be tamed in a similar

[35]Two representations $(\pi_1, \mathcal{H}_1)$, $(\pi_2, \mathcal{H}_2)$ are unitarily equivalent if there exists a unitary operator $V : \mathcal{H}_1 \to \mathcal{H}_2$ such that $V\pi_1(A) = \pi_2(A)V$ for every $A \in \mathfrak{A}$. Such representations agree on $\mathfrak{A}$, as well as on all representation-dependent parochial observables. The also have the same folium of states (physically salient states which can be represented by density operators). Unitary equivalence is generally taken as a sufficient condition for physical equivalence of representations.

fashion by a finite-volume cutoff procedure, but only in situations where the long distance physics are insensitive to the short distance physics. More typically they are dealt with by fixing asymptotic boundary conditions. This amounts to choosing different values for global observables like charge, temperature, stress-energy, net magnetization, etc. So picking different boundary conditions just is the same thing as selecting a particular representation to work in (i.e a global state). AQFT can certainly help us understand these different boundary conditions better (Wallace admits as much), but in the case of long-distance physics, the algebraic methods are actually compatible with and complimentary to the effective field theory picture. Moreover, all of the explanatory insights provided by inequivalent representations that Fraser cites concern long-distance physics.

Third, it turns out that in the Wightman framework we already have a way of recapturing the notion of an inequivalent representation while working in a single fixed Hilbert space. In AQFT, the inequivalent representations which are explanatorily significant are representations of the *observable* algebra. In the Wightman framework, we instead start with a collection of *field* operators acting on a single Hilbert space $\mathcal{H}$. If a gauge symmetry is present, the physical observables are described by the gauge invariant combinations of the field operators. The action of the gauge group also decomposes $\mathcal{H}$ into a direct sum of irreducible subspaces $\mathcal{H} = \oplus\mathcal{H}_\sigma$. The space of physically possible states is then represented by this decomposition, not the larger Hilbert space.[36] Moreover, the irreducible subspaces themselves correspond to unitarily inequivalent representations of the observable algebra. Thus once we get clear on the gauge invariant content of a constructive model, we typically find a whole host of hidden representations described by isolated sectors of a larger Hilbert space. Insofar as the Wightman formalism is continuous with the Lagrangian approach (at least for systems with global rather than local gauge symmetry), we can appeal to the same idea to make sense of inequivalent representations in the Lagrangian setting.

[36]Vectors that do not lie within an irreducible subspace are indistinguishable from mixed states with respect to the observables. Their distinct statistical properties can only be teased apart by looking at their expectation values for non-gauge invariant operators.

The main moral here is that the central issue in the Fraser-Wallace debate has to do with differences between exact QFTs and effective QFTs with a short distance cutoff. But this issue is completely orthogonal to the relationship between Lagrangian, constructive, and algebraic QFT. It should also be emphasized that even if we think the Standard Model is an effective theory, it is still worth studying exact QFTs which do not employ a cutoff. The caveat here is that we should be wary about any metaphysical conclusions drawn from such investigations, at least as they pertain to the actual world. Classical mechanics provides an instructive example in this case. Viewed as a limit of special relativity, it supports a very different metaphysical interpretation than it does viewed as an independent, exact theory in its own right.

1.3 Towards a More Cosmopolitan Worldview

If we don't follow Wallace and Fraser's lead, choose a side, and stick to our guns, what should we as philosophers of physics do? I think the answer is fairly obvious: we should approach this complex situation in much the same way that working physicists and mathematicians do.

At a recent foundational conference at Stony Brook, the famous string theorist Ed Witten remarked, "Personally I think the relationship between mathematics and physics will remain unsatisfactory unless the program of constructive field theory is reinvigorated in some form."[37] In the talk that followed, Witten sought to draw connections between issues in mainstream particle physics and open mathematical questions in constructive and algebraic field theory. This sentiment is echoed by Rudolf Haag, one of the founding fathers of AQFT. Commenting on the relationship between axiomatic and Lagrangian field theory at the end of his book *Local Quantum Physics*, he writes "It seems to me unwise to limit

[37] Simons Center for Geometry and Physics at Stony Brook, *Mathematical Foundations of Quantum Field Theory*, January 2012.

attention to one of the two approaches. [...] We need a synthesis of the knowledge gained in the different approaches."[38]

As we have seen in the course of our brief tour, each approach has different strengths and weaknesses. Lagrangian field theory supplies the raw data, a set of tremendously powerful, if somewhat piecemeal, predictive and explanatory schemas which are central to our understanding of nature at its most fundamental level (even if they're a bit rough around the edges). Everyone's goal is to understand these schemas better, and they put significant constraints on our theorizing. Axiomatic approaches give us more precise regimentations of Lagrangian QFT, but we're still in the process of figuring out how they work and if they're fully faithful to the original Lagrangian picture. Constructive field theory has numerous sophisticated tools for building concrete models of QFT. Its main drawbacks are its reliance on localized gauge-dependent field operators. AQFT stands to provide a more physically transparent, gauge free description of QFT, but the construction of particular models is very challenging. In addition, we still lack the resources to talk about certain central physical ideas in the algebraic framework. Therefore we need to look to the other, more concrete approaches for clues as to how to bring algebraic methods to bear on a wider range of problems, as well as for ideas about possible supplements or extensions of the Haag-Kastler axioms. Arguably the best way to advance this goal is to approach the problem from a number of different, complementary angles.

The result is a methodological viewpoint akin to political cosmopolitanism. It's a view on which each of these ostensibly rival programs are in fact part of a larger research culture, unified by a set of common goals and methodology. At the same time, the differences between the approaches are as significant as their similarities, highlighting different facets of QFT and drawing attention to gaps in our understanding which require further study. The path forward requires the development of each viewpoint to the fullest extent, as well as

[38]Haag (1996), p. 326. This isn't to say that there is universal harmony amongst physicists working on these various research programs. Quite the opposite. But is does seem that as a group, physicists are more open to the possibility of progress coming from different directions.

significant crosstalk between them. Admittedly, insofar as we're working towards an eventual unification of these different formalisms, our cosmopolitanism is instrumental rather than an end in itself. Still part of the cosmopolitan spirit requires an openness to the possibility that no single unified framework will be found. That would be a truly interesting discovery.

To sum things up: we have three complementary, yet partial pictures of QFT, pictures with significant overlap, differential advantages, and no manifest incompatibility. At this stage it would be premature to cast any of them aside. The moral for philosophers of QFT is this: we should be more open to ideas, puzzles, and techniques from all three approaches. To date, philosophers have focused on a fairly narrow set of questions in the foundations of QFT, many of them motivated by particular quirks in the algebraic formalism. In the meantime, there are a host of overlooked conceptual questions motivated by central physical problems arising in Lagrangian and constructive QFT which mathematically-oriented philosophers of physics are well-situated to make progress on. Some of these include the infrared problem in QED, the invariant structure of theories with local gauge symmetries, renormalization, the existence of a mass-gap, modifications to the standard scattering picture, and anomalous symmetry breaking. Axiomatic approaches to field theory are still evolving, and if philosophers wish to actively participate in their development, we must adopt a more cosmopolitan attitude towards the various conceptual schemas employed by working physicists.

My project over the next several chapters can be viewed as an initial sortie into this territory. Its main subject, *Tomita-Takesaki modular theory*, represents one of the more technical and under-explored corners in the foundations of QFT. Although modular theory dates back to the 1960s and is a mainstay of mathematical work on operator algebras, its full range of physical applications is just becoming apparent. Physicists and philosophers alike are largely unaware of the central role it plays in QFT. At a very general level, modular theory studies certain reflection symmetries of the local observable algebras encoded by a unique pair of state-dependent modular operators, (J, Δ). This pair encodes vital information about the internal algebraic structure of the local algebras, the background state, as well as the

relational structure of the net. In addition, there are a number of surprising connections between modular theory and spacetime geometry.

An important unsung theorem, the *Bisognano-Wichmann theorem*, provides a direct link between (J, Δ) and spacetime symmetries in certain physical contexts. In Ch. 3, I examine how this link can be used to give new geometric insight into the PCT theorem. I argue that modular theory is the key for unlocking the mysterious connection between spatiotemporal orientation and charge structure entailed by the theorem. In Ch. 4, I turn to the problem of localization in QFT. Modular theory tells us that at a fundamental level, the world is more thoroughly entangled than we could have possibly imagined. In spite of this, I contend that modular theory also supplies critical tools needed to help explain approximately-localized emergent structure like particles and chairs. In Ch. 5, I explore a more speculative application of the Bisognano-Wichmann theorem to solve the problem of time in quantum gravity.

Throughout these investigations, we find that modular theory sits at the nexus between Lagrangian, constructive, and algebraic QFT. While mathematically grounded in AQFT, the explanatory utility of modular theory comes from its connections to physical ideas and conceptual puzzles framed in standard Lagrangian QFT. At the same time, our current understanding of modular theory, especially its relationship to spacetime symmetries, comes from the detailed study of modular operators in particular constructive models. It is only by appreciating how these three approaches fit together then, that we can begin to piece together the full story of the physics.

From a mathematical standpoint, modular theory works by establishing a translation manual between the algebraic structure of physical quantities and the geometric structure of Hilbert space. Insofar as Lagrangian and constructive field theory hew more closely to traditional Hilbert space formulations of quantum mechanics, we should expect modular theory to give us clues about the relationship between them and the algebraic approach. This intuition is born out by a number of specific applications:

- We have seen that certain energy-bounds yield sufficient conditions for translating between constructive and algebraic QFT. Similar bounds are also needed in Lagrangian QFT to secure a coherent particle interpretation. The modular operator, Δ, encodes a number of useful facts about the local energy. In fact, generalized versions of these bounds can be framed using the condition of *modular nuclearity* (see Ch. 4.1).

- Modular theory guarantees that every local observable algebra has a canonical group of strongly-continuous automorphisms, the *modular automorphism group*. This group possesses elegant analyticity properties which are crucial surrogates for the analyticity features of Wightman fields in proofs of global structure theorems like the PCT and spin-statistics theorems.

- Modular constructions have produced new model-building techniques for AQFT, yielding some of the first examples of rigorous interacting theories in $d = 4$ spacetime.[39]

- Geometric properties of (J, Δ) can be used to extend algebraic and constructive axioms to curved spacetime settings, thereby connecting these approaches to more traditional Lagrangian methods, as well as providing a stepping stone towards an eventual generally covariant theory of quantum gravity.

While its mathematical utility has been firmly established, and we already have several examples from QFT where modular techniques are explanatorily illuminating, the broader physical significance of modular theory remains murky. This cluster of problems represents an area where I think philosophers could make very important contributions to an ongoing research program in mathematical physics.

To go any further, we first need to understand the mathematics of modular theory in more detail. We turn to this task in Ch. 2. Subsequent chapters will make use of a number of technical concepts from AQFT introduced in Appendices A and B. Experts may wish to

[39] See Summers (2012a) for a review of these results.

skip over these sections, but readers new to the material are encouraged to consult them before pressing forward.

I have argued that in order to make progress on some truly thorny foundational questions, we need more exchange between Lagrangian, constructive, and algebraic QFT, and more collaboration between researchers in physics, philosophy, and mathematics. Philosophers in particular should be more bold and look towards current physics for interesting conceptual problems and new tools to experiment with. Results from modular theory look like they could provide an intriguing catalyst as we work towards this goal.

Appendix A: Operator Algebras

Because of the failure of the Stone-von Neumann uniqueness theorem for systems with infinitely many degrees of freedom, AQFT axiomatizes algebraic collections of observables rather than quantum statespace itself. The mathematical theory of operator algebras is central to the theory.[40] It is typically assumed that the observables of a quantum system form (the self-adjoint part of) a noncommutative C^*-algebra.[41] Any such algebra is structurally equivalent to a subalgebra of bounded linear operators on some Hilbert space. We arrive at a formal definition by abstracting away from the details of the underlying space:

C^*-algberas. a *C^*-algebra*, $\mathfrak{A}$, is a complex algebra — a linear vector space over $\mathbb{C}$ with a multiplication operation (denoted AB for $A, B \in \mathfrak{A}$, with multiplicative unit I). In quantum theories, the C^*-product is noncommutative. $\mathfrak{A}$ is also required to be closed under a canonical

[40] See Kadison and Ringrose (1997a,b) for a thorough introduction to the subject.

[41] Clifton et al. (2003) argue that the operator algebra framework is broad enough to capture both classical and quantum theories. Within the space of C^*-algebraic systems, quantum theories are characterized by a unique combination of non-commutativity, non-locality, and kinematic independence. Every unital, commutative C^*-algebra is isomorphic to $C(X)$, the set of continuous, complex-valued functions on a compact Hausdorff space, X. The space X is itself isomorphic to the pure statespace over the relevant algebra endowed with the ultraweak topology. Due to this representation theorem, Clifton *et al.* conclude that commutative C^*-algebra systems are in 1-1 correspondence with classical theories. Non-commutativity is therefore a necessary condition for quantum theories.

involution mapping $* : \mathfrak{A} \to \mathfrak{A}$, satisfying:

$$(A^*)^* = A, \quad (A+B)^* = A^* + B^* \tag{1.1}$$
$$(cA)^* = \bar{c}A^*, \quad (AB)^* = B^*A^*,$$

where $c \in \mathbb{C}$, and $\bar{c}$ denotes complex conjugation. Additionally, $\mathfrak{A}$ must be complete with respect to a norm $\| \circ \|$, such that,

$$||A^*A|| = ||A||^2, \quad ||AB|| \leq ||A||\,||B||. \tag{1.2}$$

The topology induced by the norm is called the *uniform topology*. A sequence of operators A_n converges uniformly to A iff $||A_n - A|| \to 0$ as $n \to \infty$.

The involution is the abstract analogue of the adjoint operation (i.e. conjugate transpose in matrix algebras), and allows us to define important subclasses of elements in $\mathfrak{A}$. An operator, A, is said to be

self-adjoint if $A = A^*$,

positive if $A = H^*H$ for some self-adjoint $H \in \mathfrak{A}$,

a *projection* if $A^*A = A$,

isometric if $A^*A = I$,

unitary if $A^*A = AA^* = I$.

In quantum theories, observables are almost always represented by self-adjoint operators (these are the abstract analogues of hermitian operators in standard QM). Unitary operators are linked to symmetries by Wigner's theorem. Projection operators are important in the analysis of quantum measurement and quantum logic.

The *spectrum* of an operator $\mathrm{sp}(A)$ is the set of complex numbers λ such that $A - \lambda I$ does not have a two-sided inverse (i.e. an element $B \in \mathfrak{A}$ such that $B(A-\lambda I) = (A-\lambda I)B = I$).

This generalizes the notion of an eigenvalue from standard QM to cover cases of operators with continuous spectra. If $\mathfrak{A}$ is a concrete C^*-algebra (i.e. a subalgebra of $B(\mathcal{H})$) and there exists a vector $x \in \mathcal{H}$ such that $(A - \lambda I)x = 0$, then $\lambda \in \mathrm{sp}(A)$. The converse is not true in general.

Representations. A representation is a mapping from an abstract operator algebra into a subalgebra of bounded linear operators on a Hilbert space that preserves all algebraic structure. Formally, a *representation* of $\mathfrak{A}$ is a pair, $(\pi, \mathcal{H})$, where $\mathcal{H}$ is a Hilbert space and π is a *-homomorphism into $B(\mathcal{H})$, i.e. a linear mapping such that $\pi(AB) = \pi(A)\pi(B)$, $\pi(I) = I$, and $\pi(A^*) = \pi(A)^*$. If π is 1-1, the representation is said to be *faithful*. A vector $x \in \mathcal{H}$ is said to be *cyclic* for $\pi(\mathfrak{A})$ if $\pi(\mathfrak{A})x$ is dense in $\mathcal{H}$, i.e. the closed linear span $\overline{\pi(\mathfrak{A})x} = \mathcal{H}$. A vector is *separating* for $\pi(\mathfrak{A})$ if $\pi(A)x = 0$ entails that $\pi(A) = 0$.

A representation $(\varphi, \mathcal{K})$ is a *subrepresentation* of $(\pi, \mathcal{H})$ if there exists an isometry $V : \mathcal{K} \to \mathcal{H}$ such that $V\varphi(A) = \pi(A)V$. A representation is *irreducible* if it has no non-trivial subrepresentations.

Two representations $(\pi_1, \mathcal{H}_1)$ and $(\pi_2, \mathcal{H}_2)$ are *unitarily equivalent* if there exists a unitary operator $U : \mathcal{H}_1 \to \mathcal{H}_2$ such that $U\pi_1(A) = \pi_2(A)U$ for all $A \in \mathfrak{A}$. Two representations are *quasiequivalent* if they have unitarily equivalent subrepresentations. Unitary equivalence/quasiequivalence are frequently taken as sufficient conditions for physical equivalence in quantum theories.

Within a particular representation, we can define additional topologies on $B(\mathcal{H})$ and thus on $\pi(\mathfrak{A})$. Three of the most important are the *strong*, *weak*, and *ultraweak* topologies. A sequence of operators, A_n, converges to A

strongly if $|(A_n - A)|\psi\rangle| \to 0$ as $n \to \infty$ for every $|\psi\rangle \in \mathcal{H}$,

weakly if $|\langle\phi|(A_n - A)|\psi\rangle| \to 0$ as $n \to \infty$ for every $|\psi\rangle, |\phi\rangle \in \mathcal{H}$,

ultraweakly if $|\mathrm{Tr}(\rho(A_n - A))| \to 0$ as $n \to \infty$, for every density operator $\rho \in \mathcal{H}$.

In all of the above, $|\circ|$ represents the vector norm on $\mathcal{H}$. As the names suggest, the topologies are ordered $weak \subseteq strong \subseteq uniform$, and $weak \subseteq ultraweak \subseteq uniform$. (The ultraweak and strong topologies are incomparable) All four topologies coincide for finite dimensional Hilbert spaces.

States. A *state* on a C^*-algebra is a linear functional $\phi : \mathfrak{A} \to \mathbb{C}$, such that $\phi(A^*A) \geq 0$ for all $A \in \mathfrak{A}$ and $\phi(I) = 1$. A state is said to be

faithful if $\phi(A) = 0$ implies $A = 0$ for all $A \in \mathfrak{A}$,

tracial if $\phi(AB) = \phi(BA)$ for all $A, B \in \mathfrak{A}$,

mixed if $\phi = c\phi_1 + (1-c)\phi_2$, with states $\phi_1 \neq \phi_2$, $c \in (0,1)$,

pure if it is not mixed.

Note that in a given representation, vector states are not always pure in contrast to standard QM.

Given a state, ϕ, the *Gelfand-Naimark-Segal (GNS) construction* ensures the existence of a privileged home representation $(\pi_\phi, \mathcal{H}_\phi)$ associated with the state. In the GNS representation, the state ϕ is represented by a vector Φ, such that $\phi(A) = \langle \Phi, A\Phi \rangle$ for all $A \in \mathfrak{A}$. Moreover, Φ is cyclic for $\pi(\mathfrak{A})$. If ϕ is faithful, then Φ is also separating.

As we'll see in chapter 4, the local algebras of observables in AQFT typically do not have any pure states. Nonetheless in the local GNS representation, every state is a vector state.

von Neumann Algebras. The internal structure of the local algebras needs to be rich enough to capture the full range of functional relations between physical quantities localized in that region. A C^*-algebra has nearly everything we need, but there are a number of quantities of physical importance that cannot be constructed with the tools on the table thus far. Some examples include temperature, the stress-energy tensor, net magnetization of an infinite spin chain, and (in Fock space theories) the global number operator. In addition,

C^*-algebras lack a complete set of projection operators and a well-defined pair of modular operators.

To remedy these problems, we require that the local algebras be a special kind of C^*-algebra called a *von Neumann algebra*. Typically, a von Neumann algebra, $\mathfrak{M}$, is defined concretely by choosing a specific representation of a C^*-algebra and taking its closure in the weak-topology $\mathfrak{M} := \pi(\mathfrak{A})^-$. The *commutant* of a von Neumann algebra is defined as $\mathfrak{M}' := \{B \in B(\mathcal{H}) \mid [A, B] = 0,\ \forall A \in \mathfrak{M}\}$. Von Neumann's famous double-commutant theorem shows that $\pi(\mathfrak{A})^- = \pi(\mathfrak{A})''$, providing an equivalent algebraic definition. A von Neumann algebra can also be defined abstractly as a C^*-algebra dual to a Banach space. (Some authors refer to the abstract definition as a W^*-*algebra*.) Throughout, we will use the pair of equivalent concrete definitions.[42]

The *center* of a von Neumann algebra is defined as $Z(\mathfrak{M}) := \mathfrak{M} \cap \mathfrak{M}'$. A *factor* is a von Neumann algebra with a trivial center, $Z(\mathfrak{M}) = \mathbb{C}I$. In this case $\mathfrak{M} \vee \mathfrak{M}' = B(\mathcal{H})$, i.e. $B(\mathcal{H})$ factorizes into $\mathfrak{M}$ and its commutant. Throughout, we'll use $\mathfrak{M} \vee \mathfrak{N} := (\mathfrak{M} \cup \mathfrak{N})''$ to denote the joint von Neumann algebra generated by a pair of von Neumann algebras.

A *normal state* on a von Neumann algebra is an ultraweakly continuous state. Such states are countably additive in the sense that $\sum_i \phi(E_i) = \phi(\sum_i E_i)$, for countable families of projection operators $\{E_i\}$. Moreover, normal states stand in 1-1 correspondence with density operators on $B(\mathcal{H})$. It is typically assumed that all states of direct physical interest

[42]Given the difficulties associated with inequivalent representations, we adopt the following perspective regarding the origins of the local von Neumann algebras. Given a net of C^*-algebras satisfying the Haag-Kastler axioms, the net of local von Neumann algebras is defined by taking the GNS representation associated with the restriction of the global state to each region of spacetime. If we impose a set of principled, physically motivated restrictions on possible global states (e.g. the DHR/BF condition), this latter strategy corresponds to the interpretive stance Ruetsche (2011a) dubs *representational realism*. On this view, the corresponding category of physical representations encapsulates a significant portion of the theory. It's worth noting that not only does this stance seem to align best with the views of the mathematical physics community, but it also offers the clearest way to understand the algebraic PCT and spin-statistics theorems. From this angle, the arguments in Ch. 3 can be seen as another salvo in the philosophical debate over the right way to interpret AQFT. Why be a representational realist? Because it offers the best explanation of the PCT theorem and spin-statistics connection. Rigorous proofs of these theorems are frequently cited as among the most important successes of AQFT.

are normal. See Ruetsche (2011b) for a number of possible explanations. The collection of all normal states for a von Neumann algebra $\mathfrak{M}$ is called the *folium* of $\mathfrak{M}$.

Type Classification. The Murray-von Neumann type classification of von Neumann algebras is based on the structure of their lattice of projections. A projection $F \in \mathfrak{M}$ is a *subprojection* of $E \in \mathfrak{M}$ if $F\mathcal{H} \subset E\mathcal{H}$. Two projections $E, F \in \mathfrak{M}$, are *equivalent* (modulo $\mathfrak{M}$), written $E \sim F$, if there exists a partial isometry $V \in \mathfrak{M}$ such that $V^*V = E$ and $VV^* = F$. (A partial isometry is a linear map between Hilbert spaces whose restriction to the orthogonal complement of its kernel is an isometry. Examples include unitaries and projections.) A projection operator, $E \in \mathfrak{M}$, is:

minimal when $\mathfrak{M}$ contains no proper subprojection $F < E$,

abelian when the von Neumann algebra $E\mathfrak{M}E$ is abelian,

infinite when $\exists F \in \mathfrak{M}$ such that $E \sim F < E$,

finite when E is not infinite,

finite-dimensional when the range of E is a finite-dimensional subspace of $\mathcal{H}$.

In general, the following implications hold:

$$\text{finite-dimensional} \Rightarrow \text{minimal} \Rightarrow \text{abelian} \Rightarrow \text{finite} \Leftrightarrow \neg\text{infinite}$$

For factors, a projection is minimal iff it is abelian. A factor is then said to be

Type I if it contains a non-zero abelian (equiv. minimal) projection,

Type II if it contains a non-zero finite projection, but no abelian projections,

Type III if every projection is infinite.

Cutting across this type classification, a factor is said to be of *finite type* if the identity projection, I, is a finite projection (type I_n, II_1), and of *infinite type* if I is an infinite

projection (type I_∞, II_∞, III). Every type I_n factor is isomorphic to $B(\mathcal{H}_n)$ for a Hilbert space $\mathcal{H}_n$ with finite dimension n. Infinite factors can only be concretely implemented on an infinite dimensional Hilbert space.

Appendix B: The Haag-Kastler Axioms

The Haag-Kastler axioms attempt to capture the constraints that relativity and QM imposes on the collection of local algebras. The first three axioms define net of C^*-algebras over Minkowski spacetime and ensure that they are properly local. Throughout we will let $\mathcal{O}$ denote any open, bounded region of spacetime, and $\mathcal{D}$ a double cone formed by the intersection of a future and past lightcone. $\mathfrak{A}(\mathcal{O})$ and $\mathfrak{R}(\mathcal{O}) := \pi(\mathfrak{A}(\mathcal{O}))''$ will denote local C^*- and von Neumann algebras respectively, while $\mathfrak{A}$ and $\mathfrak{R} := \mathfrak{A}''$ will denote the quasilocal and global algebras.

Isotony (1) If $\mathcal{O}_1 \subseteq \mathcal{O}_2$, then $\mathfrak{A}(\mathcal{O}_1)$ is a subalgebra of $\mathfrak{A}(\mathcal{O}_2)$. The quasi-local algebra $\mathfrak{A}$ is defined as the norm-closure of $\bigcup \mathfrak{A}(\mathcal{O})$ over all regions $\mathcal{O}$.

Microcausality (2) If $\mathcal{O}_1 \subseteq \mathcal{O}_2$ are spacelike separated, then $[\mathfrak{A}(\mathcal{O}_1), \mathfrak{A}(\mathcal{O}_2)] = \{0\}$.

Another locality axiom, microcausality captures the idea that spacelike separated regions form causally independent subsystems.[43] If there were some pair of observables $A \in \mathfrak{A}(\mathcal{O}_1)$, $B \in \mathfrak{A}(\mathcal{O}_2)$, that did not commute, then a measurement of one could effect the statistics of the other at spacelike separation in violation of relativistic causality.

Covariance (3) The net of local algebras is covariant under the action of continuous Poincaré transformations, i.e. there exists a faithful, continuous representation of $\mathcal{P}_+^\uparrow$ in $\mathrm{Aut}(\mathfrak{A})$ such that for any region, $\mathcal{O}$, and any $g \in \mathcal{P}_+^\uparrow$,

$$\alpha_g(\mathfrak{A}(\mathcal{O})) = \mathfrak{A}(g\mathcal{O}). \tag{1.3}$$

[43] It should be noted that there are a number of different criteria which might be thought of as the right notion of causal independence. For a survey of recent results see Halvorson Muger §3.

It is important to note that the locality assumptions listed above do not conflict with Bell's Theorem. In a somewhat unfortunate use of terminology, physicists refer to local quantum field theories as those whose physical *quantities* satisfy locality constraints like those listed above. The notion of locality relevant for of Bell's Theorem concerns *states*. Thus in a local quantum field theory, observables are localized, but states can be as non-local as you like. In fact a number of results suggest that maximally non-local states are generic in quantum field theory unlike in ordinary QM (see Ch. 4).

The next two axioms concern special properties of the vacuum state.

Vacuum (4) There exists a unique state, ω, invariant under under the translation subgroup $\mathcal{T} \subset \mathcal{P}_+^\uparrow$.

As a consequence of (4), in the vacuum GNS representation $(\pi_\omega, \mathcal{H}_\omega)$, the translations are implemented by unitary transformations, i.e. $\alpha_{g(a)}(\mathfrak{A}) = U(a)\mathfrak{A}U(-a)$ where $U(a)$ is a unitary element of $\mathfrak{A}$. Since $\mathcal{T}$ is a locally compact abelian topological group, we can use the Stone-Naimark-Ambrose-Godement (SNAG) theorem to write $U(a)$ as an integral over momentum space:

$$U(a) = \int_{\hat{\mathcal{T}}} p(a)\ d\mu(p). \tag{1.4}$$

Here $\hat{\mathcal{T}}$ is the dual of $\mathcal{T}$. Its elements are group homomorphisms $f : \mathcal{T} \to \mathbb{C}/\{0\}$, the characters of $\mathcal{T}$. If $a \in \mathcal{T}$ represent translation 4-vectors, $p \in \hat{\mathcal{T}}$ represent dual energy-momentum 4-vectors. The measure, μ, is the unique regular spectral measure on $\hat{\mathcal{T}}$ whose existence is guaranteed by the theorem. It takes Borel sets $I \subset \hat{\mathcal{T}}$ to associated projection operators $\mu(I)$ acting on H_ω.[44]

As usually formulated the *spectrum condition* requires the support of μ to lie in the forward lightcone in momentum space. The apparent reference to a preferred temporal direction is inessential:

[44]See Folland (1995, Thm. 4.44) for a proof.

Spectrum Condition (5) In the vacuum GNS representation, supp(μ) must lie in a closed, convex set $\overline{V} \in \hat{\mathcal{T}}$ which is asymmetric under taking additive inverses, i.e. $\{\overline{V}\} \cap \{-\overline{V}\}=\{0\}$.

The standard Haag-Kastler axioms are rounded out by one additional technical assumption:

Weak Additivity (6) For any double cone $\mathcal{D}$, the global algebra $\mathfrak{R} := \pi(\mathfrak{A})''$ can be generated by translations of any local algebra, i.e. $\mathfrak{R} = \bigvee \mathfrak{R}(\mathcal{D} + a)$ over all translations $a \in \mathcal{P}_+^\uparrow$.

The standard justification for this axiom is the requirement that there should be no smallest length scale in the theory; the local algebras are non-trivial for regions of arbitrary size. This assumption could fail if we view the theory as an effective field theory approximating some other underlying theory (e.g. string theory, quantum gravity).

Many physically interesting models satisfy further constraints. Occasionally, we will appeal to the following additional axioms:

Strong Additivity For any family of regions $\{\mathcal{O}_i\}$, $\mathfrak{R}(\cup\mathcal{O}_i) = \bigvee \mathfrak{R}(\mathcal{O}_i)$.

Strong additivity amounts to the requirement that all the observables localized in region $\mathcal{O}$ can be constructed out of observables from any open cover. Although it is known to fail in certain situations, strong additivity holds in a wide range of physically interesting cases.[45] It entails both isotony and weak additivity.

Along with strong additivity, two more conditions are required to ensure that the local algebras, $\mathfrak{R}(\mathcal{O})$, are factors:

[45]For example, given a net of field operators, $\mathfrak{F}$, satisfying strong additivity, Roberts (1982, Thm. 1.3) shows that if $\mathfrak{R}(\mathcal{O}) = \{A \in \mathfrak{F}(\mathcal{O}) \mid gA = A, g \in G\}$ for some compact gauge group G, then the observable net satisfies strong additivity. Strong additivity for the field net, $\mathfrak{F}$, in turn, is supported by observations from Wightman field theory. Given a partition of unity in the relevant test function space, if $\mathcal{O} = \cup\mathcal{O}_i$ and f is a test function with support in $\mathcal{O}$ we can write $f = \sum f_i$ where f_i have support in $\mathcal{O}_i$. For any Wightman field, it follows that $\psi(f) = \sum \psi(f_i)$, and thus polynomials of the fields must satisfy strong additivity. Generalizing this result requires dealing with thorny questions concerning whether or not the net of unbounded field operators is essential self-adjoint. In many concrete models these difficulties can be overcome. For example free-Bose and free-Fermi theories satisfy strong additivity, as well as $P(\phi)_2$ theories and QFTs generated by Euclidean Green's functions.

Inner/Outer Continuity A net of von Neumann algebras is *inner continuous* if for any monotonically increasing net $(\mathcal{O}_a)$ with least upper bound $\mathcal{O}$, $\bigvee \mathfrak{R}(\mathcal{O}_a) = \mathfrak{R}(\mathcal{O})$. The net is *outer continuous* if for any decreasing net $(\mathcal{O}_b)$ with greatest lower bound $\mathcal{O}$, $\bigcap \mathfrak{R}(\mathcal{O}_b) = \mathfrak{R}(\mathcal{O})$.

Primitive Causality For any *complete spacelike slice*, $\mathcal{S}$, $\mathfrak{R}(\mathcal{S}) = \mathfrak{R}$.

Intuitively, a complete spacelike slice is a "thickened" cauchy surface.[46] Primitive causality (also called the *time-slice axiom*), ensures the existence of a globally hyperbolic set of field equations (even if these equations are too complex to express in closed form).

One final condition which will appear often is *duality*:

Duality For any double cone $\mathcal{D}$ with spacelike complement $\mathcal{D}'$, $\mathfrak{R}(\mathcal{D}') = \mathfrak{R}(\mathcal{D})'$.

Duality says that the local double-cone algebras are maximal. It is only expected to hold in certain restricted cases (see the discussion on p. 141). For certain unbounded regions including spacelike wedges, duality is expected to hold quite broadly. As we'll see in Ch. 2, wedge-duality is a direct consequence of the Bisognano-Wichmann theorem.

Appendix C: The Cutoff Reeh-Schlieder Theorem

Here we formulate and prove a rigorous version of Prop. 1.1.

A vector, $\Phi \in \mathcal{H}$ is *analytic for the energy*, if the series

$$\sum_{n=0}^{\infty} ||(P^0)^n \Phi|| \, \frac{z^n}{n!} \tag{1.5}$$

is convergent in any circle of finite radius in the complex plane. It follows that for any unitary operator implementing a translation $U(a)$ the vector-valued function $U(a)\Phi$ has an analytic continuation in a to a function holomorphic on $\mathbb{C}^4$. The vacuum state vector, Ω, is

[46] We can define the slice $\mathcal{S}_\epsilon^{(n)} := \{x \in \mathcal{M} \mid |n_a x^a| < \epsilon\}$, for some timelike normal 4-vector n_a.

analytic for the energy, as are vectors representing states of compact energy (which include all states satisfying the DHR/BF selection criteria).

Given two spacetime regions, we say that $\mathcal{O}_1$ is *strictly contained* in $\mathcal{O}_2$, denoted $\mathcal{O}_1 \ll \mathcal{O}_2$, if there exists a neighborhood, $\mathcal{N}$, of the origin such that for all $n \in \mathcal{N}$, $\mathcal{O}_1 + n \subset \mathcal{O}_2$ and there exists $\mathcal{O}_3 \subset \mathcal{O}_2$, such that $\mathcal{O}_1$ and $\mathcal{O}_3$ are spacelike separated.

Assuming $\lambda \in \mathbb{R}^+$ is a positive scale factor, let $\mathcal{D}_\lambda$ denote a double-cone with cutoff volume λ. We can formulate a cutoff-safe version of weak additivity as follows:

Cutoff Weak Additivity (6*) For any double cone $\mathcal{D}$ such that $\mathcal{D}_\lambda \ll \mathcal{D}$, the global von Neumann algebra $\mathfrak{R} = \bigvee \mathfrak{R}(\mathcal{D} + a)$ over all translations $a \in \mathcal{P}_+^\uparrow$.

With these definitions in place can state our main lemma:

Proposition 1.1 (Cutoff Reeh-Schlieder)**.** *Let Φ be a vector analytic for the energy. If the specrum condition and cutoff weak additivity hold, then $\overline{\mathfrak{R}(\mathcal{D})\Phi} = \mathcal{H}$ for all double-cones $\mathcal{D} \gg \mathcal{D}_\lambda$. If Microcausality holds, then Φ is separating for $\mathfrak{R}(\mathcal{D})$.*

Proof. Since $\mathcal{D} \gg \mathcal{D}_\lambda$, we can find an interpolating double-cone $\tilde{\mathcal{D}}$ such that $\mathcal{D} \gg \tilde{\mathcal{D}} \gg \mathcal{D}_\lambda$. Consequently, there exists a neighborhood of the origin such that for all $n \in \mathcal{N}$, $\mathfrak{R}(\tilde{\mathcal{D}} + n) \subset \mathfrak{R}(\mathcal{D})$. Letting V^+ denote the forward lightcone and $U(a)$ an unitary operator implementing a translation, the analyticity of Φ in conjunction with the spectrum condition ensure that the vector-valued function $U(a)\Phi$ has an analytic extension to the forward tube $T(V^+) := \{z \in \mathbb{C}^4 \mid \text{Im } z \in \overline{V^+}\}$. We can now apply the standard argument:

Let $\Psi \in B(\mathcal{H})$ be any vector orthogonal to $\mathfrak{R}(\mathcal{D})\Phi$. We will show that $\Psi = 0$. By the analyticity properties of operators of $U(a)\Phi$ discussed above, for any family of operators $A_j \in \mathfrak{R}(\tilde{\mathcal{D}})$, with $n_j \in \mathcal{N}$, the function,

$$\langle \Psi, U(n_1)A_1U(n_2 - n_1)A_2 \dots U(n_j - n_{j-1})A_jU(-n_j)\Phi \rangle, \tag{1.6}$$

is the boundary value of an analytic function

$$\langle \Psi, U(z_1)A_1U(z_2-z_1)A_2\dots U(z_j-z_{j-1})A_jU(-z_j)\Phi\rangle, \tag{1.7}$$

which is holomorphic for $z_1, z_2-z_1, \dots, z_j \in T(V^+)$. But since $U(n)AU(-n) \in \mathfrak{R}(\mathcal{D})$ for any $n \in \mathcal{N}$, (1.6) and thus (1.7) must vanish. But since (1.7) is holomorphic, if it vanishes on any neighborhood, it must vanish everywhere. Thus,

$$\langle \Psi, U(a_1)A_1U(a_2-a_1)A_2\dots U(a_j-a_{j-1})A_jU(-a_j)\Phi\rangle = 0, \tag{1.8}$$

for all translations $a \in \mathcal{P}_+^\uparrow$. By cutoff weak additivity $\bigvee \mathfrak{R}(\tilde{\mathcal{D}}+a) = \mathfrak{R}$. Moreover, $\overline{\mathfrak{R}\Phi} = \mathcal{H}$ by the GNS construction so $\Psi = 0$.

Φ is therefore cyclic for $\mathfrak{R}(\mathcal{D})$. To show that it is also separating, consider another double-cone $\mathcal{D}_1 \gg \mathcal{D}_\lambda$, spacelike separated from $\mathcal{D}$. By microcausality $\mathfrak{R}(\mathcal{D}_1) \subset \mathfrak{R}(\mathcal{D})'$. Thus if $A\Phi = 0$, for any $B \in \mathfrak{R}(\mathcal{D}_1)$, we have $AB\Phi = BA\Phi = 0$. But by the first part of the theorem, Φ must be cyclic for $\mathcal{D}_1$ too, so $A\overline{\mathfrak{R}(\mathcal{D}_1)\Phi} = A\mathcal{H} = 0$, therefore $A = 0$.

□

Chapter 2

Modular Theory for Philosophers

"A beautiful example of the prestabilized harmony between physics and mathematics" — R. Haag

The origins of Tomita-Takesaki modular theory lie in the structure theory of infinite von Neumann algebras. Tomita laid the groundwork for the theory in unpublished work from 1967 on the relationship between a von Neumann algebra and its commutant, proving a general commutation theorem for tensor products of von Neumann algebras — $(\mathfrak{M}_1 \otimes \mathfrak{M}_2)' = \mathfrak{M}_1' \otimes \mathfrak{M}_2'$. Almost immediately, parallels were recognized between the new mathematical tools developed by Tomita and the Kubo-Martin-Schwinger (KMS) condition characterizing equilibrium states in quantum statistical mechanics and QFT. In a subsequent monograph, Takesaki (1970) regimented the foundations of Tomita's theory and forged a rigorous mathematical connection between modular theory and the KMS condition. Thus since its inception, modular theory has occupied a position in the rich trading zone between mathematics and physics.

As it developed further, modular theory continued to provide deep insight into the mathematical foundations of QFT. It proved to be the key to understanding the structure of mysterious type III von Neumann algebras. Originally dismissed as physical pathologies by von Neumann, these were later shown to be of tremendous importance in QFT. The ground-

breaking work of Connes (1973) provided a classification scheme and structure theory for type III algebras based on modular invariants. Of nearly equal importance is the seminal theorem of Bisognano and Wichmann (1975), which showed that in certain physical settings, the local modular operators in QFT act geometrically as elements of the Poincaré group. This surprising link between modular and spatiotemporal structure lies behind many of the most important applications of modular theory in QFT.

This chapter represents a short primer on modular theory for philosophers of physics. The first part covers the standard mathematical results at the heart of the theory, emphasizing the role that modular theory plays in linking algebraic and geometric structure on Hilbert space. The second part discusses a number of important physical applications centered around the Bisognano-Wichmann theorem.[1]

2.1 The Mathematics

The mathematical setting for modular theory is a von Neumann algebra, $\mathfrak{M}$, acting on a separable Hilbert space, $\mathcal{H}$, with a cyclic, separating vector, Ω. The *spatial theory* of von Neumann algebras concerns the relationship between algebraic structure on $\mathfrak{M}$ and geometric structure on $\mathcal{H}$. Since the closed linear span of $\mathfrak{M}\Omega$ is the entire Hilbert space and the vector Ω also separates algebraic elements ($A\Omega = B\Omega \;\Rightarrow\; A = B$), we can use it to translate algebraic structure from $\mathfrak{M}$ into geometric structure on $\mathcal{H}$. In general, however, the canonical algebraic involution, *, does not give rise to an isometry of the Hilbert space structure, putting a kink in the translation. Tomita's insight was to recognize that the involution can always be represented as a reflection with an additional "twist." This served as the jumping off point for all future developments in modular theory.

[1]Material from this chapter was originally presented at a meeting of the philosophy of physics working group at the University of Michigan in January 2013. For a concise mathematical introduction to modular theory see Summers (2006). Borchers (2000) represents a more in depth survey, focusing on applications in physics. Comprehensive mathematical treatment can be found in Takesaki (2000) and Kadison and Ringrose (1997b, Ch. 9).

Using Ω, we can define the operator,

$$S_0 A\Omega = A^*\Omega \ , \tag{2.1}$$

for all $A \in \mathfrak{M}$. S_0 is unbounded and *antilinear*, i.e. $S_0(ax + by) = \bar{a}S_0 x + \bar{b}S_0 y$, $a, b \in \mathbb{C}$, $x, y \in \mathcal{H}$. Using standard techniques, we can extend S_0 to a closed, antilinear operator, S, defined on a dense subset of $\mathcal{H}$.[2] Any closed operator, A, on a Hilbert space has a unique *polar decomposition*, $A = V|A|$, where V is a partial isometry and $|A|$ is a (generally unbounded) positive, self-adjoint operator called the *modulus* of A.

In the present case, the polar decomposition of S is given by:

$$S = J\Delta^{1/2} \ , \tag{2.2}$$

with partial isometry J and modulus $(S^*S)^{1/2} = \Delta^{1/2}$. It turns out that J is an *antiunitary* operator, i.e. an antilinear operator such that $J^*J = JJ^* = I$. Furthermore, it can be shown that $J = J^* = J^{-1}$, thus J acts as a conjugation operation on $\mathcal{H}$. It is called the *modular conjugation*. The positive operator Δ is called the *modular operator*. Note that in general, any antiunitary involution on $\mathcal{H}$ implements a conjugation operation. What makes J special, a *modular* conjugation, is its relationship with Δ. In particular, $J\Delta J = \Delta^{-1}$. Both J and the spectral projections $\{E_\Delta\}$ are elements of $\mathfrak{M}$. (This is not the case for a general C^*-algebra.)

Together, the operators Δ, J have a rich structure that forms the basis of Tomita-Takesaki modular theory. The main theorem of modular theory establishes the existence of a canonical group of automorphisms of $\mathfrak{M}$ and a canonical anti-isomorphism between $\mathfrak{M}$ and its commutant $\mathfrak{M}'$:

[2]We expand the domain of definition of S_0 as follows: Let Ψ_n be a sequence in $\mathfrak{M}\Omega$ that converges strongly to $\Psi \in \mathcal{H}$. If $S_0\Psi_n$ converges strongly to some vector $\Phi \in \mathcal{H}$, then Φ is uniquely determined by Ψ. We can define $S\Phi = \Psi$. Now, let Ψ_n be a sequence that converges strongly to 0. In this case, we define the limit of $S\Psi_n = 0$.

Theorem 2.1 (Tomita-Takesaki)**.** [3] *Let $\mathfrak{M}$ be a von Neumann algebra with a cyclic and separating vector $\Omega \in \mathcal{H}$, and let Δ, J be the associated modular objects. Then*

$$(i)\ J\Omega = \Omega = \Delta\Omega\ ,$$
$$(ii)\ \Delta^{it}\mathfrak{M}\Delta^{-it} = \mathfrak{M}, \forall t \in \mathbb{R},$$
$$(iii)\ J\mathfrak{M}J = \mathfrak{M}'\ .$$

Since Δ is positive, exponentiation by it produces a unitary operator, and $\{\Delta^{it} | t \in \mathbb{R}\}$ automatically form a strongly continuous unitary group. By (ii) in the above theorem, this group induces a 1-parameter automorphism group of $\mathfrak{M}$ — the *modular automorphism group*, $\{\sigma_t\}$, defined by

$$\sigma_t(A) = \Delta^{it} A \Delta^{-it} \tag{2.3}$$

for all $A \in \mathfrak{M}$, $t \in \mathbb{R}$. By (iii), the adjoint action of the modular conjugation generates an anti-isomorphism between $\mathfrak{M}$ and $\mathfrak{M}'$.

Whenever ω is a faithful, normal state on $\mathfrak{M}$, the GNS construction yields a representation $(\pi_\omega, \mathcal{H}_\omega)$, in which ω is represented by a cyclic, separating vector $\Omega \in \mathcal{H}_\omega$. The associated modular structure provides information about both the algebra and the state. For instance, the modular automorphism group is inner iff the algebra is type I or II. It also follows directly from (i) that ω is invariant under $\{\sigma_t\}$ (i.e. $\omega(\sigma_t(A)) = \omega(A)$ for all A, t). Moreover, for any $A \in \mathfrak{M}$, $\sigma_t(A) = A$ iff $\omega(AB) = \omega(BA)$ for all $B \in \mathfrak{M}$. So ω is a tracial state iff the modular automorphism group is trivial.

An additional important fact that we will appeal to later is the following. A von Neumann algebra is called a *factor* if its center, $Z(\mathfrak{M}) := \mathfrak{M} \cap \mathfrak{M}'$, is trivial. In this case $\mathfrak{M} \vee \mathfrak{M}' = B(\mathcal{H})$, so $B(\mathcal{H})$ "factors" into the algebra and its commutant.) If $\mathfrak{M}$ is a factor (as is almost always the case for local algebras in QFT), the modular automorphism groups associated

[3] In the case where S is a bounded operator an elementary proof can be given. See Blackadar (2006, Thm. III.4.3.2.). The unbounded case is highly non-trivial. See Takesaki (2000, Ch. VI-VII) and Kadison and Ringrose (1997b, Ch. 9.2) for different versions of the full proof.

with two states ϕ and ω over $\mathfrak{M}$ are equivalent iff $\phi = \omega$. Otherwise, the groups are related by a non-trivial *inner automorphism*:

$$\{\sigma_t\}^{\phi} = U_t\{\sigma_t\}^{\omega}U_t^*, \tag{2.4}$$

where $U_t \in \mathfrak{M}$ is unitary for all t. (An automorphism is said to be *inner* if it can be implemented by unitary elements of $\mathfrak{M}$, otherwise it is *outer*.) The inner-equivalence class of modular flows associated with all faithful, normal states over $\mathfrak{M}$ corresponds to a unique, non-trivial flow of outer automorphisms of $\mathfrak{M}$. This state independent modular structure is important for classifying types of von Neuman algebras (especially the subtypes of type III).

2.1.1 The KMS Condition

In classical and quantum statistical mechanics of a large but finite number of degrees of freedom, equilibrium states are characterized by the Gibbs postulate:

$$\rho = \frac{1}{Z}e^{-\beta H} \tag{2.5}$$

where $Z = \mathrm{tr}[e^{-\beta H}]$ is the partition function and $\beta = 1/k_bT$ is the inverse temperature. Intuitively, the Gibbs postulate captures the idea that an equilibrium state is one of maximal entropy. For systems with an infinite number of degrees of freedom, the Hamiltonian operator H is unbounded and hence the partition function is not well defined. In this context, equilibrium states can be characterized using the *KMS condition*. A state, ρ, on a von Neumann algebra, $\mathfrak{M}$, is said to obey the KMS condition for inverse temperature $0 < \beta < \infty$ with respect to a 1-parameter group of automorphisms, $\{\alpha_t\}$, of $\mathcal{R}$ if for any $A, B \in R$ there exists a complex function $F_{A,B}(z)$, analytic on the strip $S(0,\beta) := \{z \in \mathbb{C} | 0 < \mathrm{Im}z < \beta\}$

and continuous on the boundary of the strip, such that

$$F_{A,B}(t) = \rho(\alpha_t(A)B) \tag{2.6}$$
$$F_{A,B}(t + i\beta) = \rho(B\alpha_t(A))$$

for all $t \in \mathbb{R}$. The KMS state ρ is invariant under the flow generated by $\{\alpha_t\}$, and is both kinematically stable under perturbations to $\{\alpha_t\}$ and passive (in the sense that a system in state ρ cannot perform net work via cyclic processes).[4] Moreover in the finite limit, the KMS condition implies the Gibbs postulate. Thus, KMS states are the appropriate generalization of the notion of an equilibrium state to infinite dimensional systems.

Takesaki (1970) establishes a remarkable connection between the KMS condition and the modular automorphism group. Given *any* faithful, normal state on $\mathfrak{M}$, the state satisfies the KMS condition with respect to the associated modular automorphism group. Here we provide a sketch of the main argument.[5]

Let ω be a faithful normal state on $\mathfrak{M}$, represented by Ω in the corresponding GNS representation. Since $A\Omega$ lies in the domain of $\Delta^{1/2}$ for any $A \in \mathfrak{M}$, the vector-valued function $t \mapsto \Delta^{it}A\Omega$ has an analytic continuation into the strip $S(0, 1/2)$. Since $SA\Omega = J\Delta^{1/2}A\Omega = A^*\Omega$, it follows that

$$\Delta^{i(t-i/2)}A\Omega = \Delta^{it}JA^*\Omega, \tag{2.7}$$

for all $A \in \mathfrak{M}$. By the definition of S, the adjoint S^* acts similarly on the commutant:

$$S^*B\Omega = B^*\Omega, \tag{2.8}$$

[4]See Haag (1996) Ch. 5.3 for more details.
[5]See Kadison and Ringrose (1997b, Thm. 9.2.13) for full details.

for all $B \in \mathfrak{M}'$. By a parallel argument, the function $\Delta^{it} B\Omega$ has an analytic continuation into the strip $S(1/2, 1)$, and thus,

$$\Delta^{i(t+i/2)} B\Omega = \Delta^{it} JB^*\Omega, \tag{2.9}$$

for all $B \in \mathfrak{M}'$. Finally, from (2.7), (2.9), and the fact that $J^2 = I$, it follows that for any $A, B \in \mathfrak{M}$, the function $t \mapsto \omega(B\sigma^t(A))$ has an analytic continuation into the strip $S(0, 1)$. At the upper boundary, we have

$$\omega(\sigma_t(A)B) = \omega(B\sigma_{t-i}(A)), \tag{2.10}$$

and thus ω satisfies the KMS condition with respect to σ_t at inverse temperature $\beta = 1$.

This result is significant for three reasons. First, it can be shown that the KMS condition characterizes the modular automorphism group uniquely. If ω satisfies the KMS condition with respect to an arbitrary group of automorphisms $\{\alpha_t\}$, then $\alpha_t = \sigma_t^\omega$ for all t. Moreover, if $\mathfrak{M}$ is a factor, then two distinct states $\omega \neq \phi$ cannot both be KMS states for the same modular flow.[6] Second, the argument indicates that the modular automorphism group comes equipped with analyticity properties associated with the KMS condition. These provide a powerful technical tool for proving theorems in AQFT where analyticity properties associated with Wightman fields are absent. Additionally, the analytic boundary condition (2.10) can be viewed as a generalization of the definition of a tracial state $\omega(AB) = \omega(BA)$. This allows us to see why tracial states correspond to trivial modular automorphism groups — whenever a state is tracial, the function $F_{A,B}$ is constant on the entire strip $S(0, 1)$. Hence, via the KMS condition, we can view the modular automorphism group as providing a rough measure of how much a state deviates from being a trace. Third, the theorem provides our first example of a direct connection between modular theory and physics, yielding an elegant

[6] More generally, a normal state ϕ satisfies the KMS condition with respect to $\{\sigma_t\}^\omega$ iff $\phi(A) = \omega(HA)$ for some positive operator H affiliated with the center of $\mathfrak{M}$.

means of characterizing equilibrium states in QFT and QSM: an equilibrium state is a state whose modular automorphism group coincides with the dynamics.

2.1.2 Connes's Classification

The classification of type III von Neumann algebras was a longstanding problem in the theory of operator algebras following the early work of Murray and von Neumann (1936). In 1972, Connes provided a solution, drawing upon the results described above. As we have already seen, if $\mathfrak{M}$ is a factor the modular automorphism groups associated with distinct states, $\phi \neq \omega$, on $\mathfrak{M}$ are related by an inner automorphism, $\{\sigma_t\}^\phi = U_t\{\sigma_t\}^\omega U_t^*$. In addition, the unitaries in $\mathfrak{M}$ implementing the inner automorphism can be chosen such that $U_{t_1+t_2} = U_{t_1}\sigma^\omega_{t_1}U_{t_2}$. The collection of unitaries $\{U_t\}$ is called the *cocycle derivative* of ω with respect to ϕ, written $U_t = (D\omega : D\phi)_t$. The cocycle derivative shares many of the same properties as an ordinary derivative and gives rise to a well-defined integration theory (via the noncommutative analog of the Radon-Nikodym theorem).

Connes used the continuity properties of the cocycle derivative to construct a canonical group homomorphism $\delta : \mathbb{R} \to \mathrm{Out}(\mathfrak{M})$, where $\mathrm{Out}(\mathfrak{M}) = \mathrm{Aut}(\mathfrak{M})/\mathrm{Inn}(\mathfrak{M})$ is the group of outer automorphisms of $\mathfrak{M}$. This homomorphism characterizes the outer flow of states on the algebra $\mathfrak{M}$ and is an invariant of the algebra.[7] Connes defined two related modular invariants,

$$\begin{aligned} T(\mathfrak{M}) &:= \ker(\delta) = \{t \in \mathbb{R} \mid \sigma_t^\omega \in \mathrm{Inn}(\mathfrak{M})\} \\ S(\mathfrak{M}) &:= \mathrm{sp}(\delta) = \{\bigcap \mathrm{sp}(\Delta_\omega) \mid \omega \text{ faitful, normal}\}, \end{aligned} \tag{2.11}$$

which leads to the following classification theorem:

[7] Connes actually worked with the broader concept of a *weight* — a positive linear functional on $\mathfrak{M}$ lacking the normalization condition for states. A von Neumann algebra possesses a faithful normal state iff it is countably decomposable, but any von Neumann algebra has a faithful, normal weight. Thus Connes's classification scheme extends to cover to all von Neumann algebras. For physical applications, it suffices to consider the more limited case of states.

Theorem 2.2 (Connes). [8] *Let $\mathfrak{M}$ be a factor on a separable Hilbert space. $\mathfrak{M}$ is type I or II iff $T(\mathfrak{M}) = 0$. If $\mathfrak{M}$ is type I_n or II_1, then $S(\mathfrak{M}) = \{1\}$. If $\mathfrak{M}$ is type I_∞ or II_∞, then $S(\mathfrak{M}) = \{0, 1\}$. If $\mathfrak{M}$ is type III, then*

(i) $\mathfrak{M}$ is type III_0 iff $S(\mathfrak{M}) = \{0, 1\}$

(ii) $\mathfrak{M}$ is type III_λ, $0 < \lambda < 1$, iff $S(\mathfrak{M}) = \{0\} \cup \{\lambda^n \mid n \in \mathbb{Z}\}$

(iii) $\mathfrak{M}$ is type III_1 iff $S(\mathfrak{M}) = [0, \infty)$

As we will see in chapter 4, the local algebras of observables in QFT are generically type III_1 factors. In the above classification, these turn out the be the most noncommutative von Neumann algebras in a well-defined sense. Connes and Størmer (1978) introduce the *state orbit diameter* $d(\mathfrak{M})$. The orbit of a state under inner automorphisms, $[\omega]$, is defined as the norm closure of $\{\omega \circ \alpha | \alpha \in \mathrm{Inn}(\mathfrak{M})\}$. The induced metric

$$d([\omega], [\phi]) = \inf\{||\omega - \phi|| \ |\omega \in [\omega], \phi \in [\phi]\} \tag{2.12}$$

can be used to measure distances between orbits. Letting $\mathfrak{M}_*$ denote the space of normal states over $\mathfrak{M}$, the state orbit diameter $d(\mathfrak{M})$ is then defined as the diameter of $\mathfrak{M}_*/\mathrm{Inn}(\mathfrak{M})$. $d(\mathfrak{M}) \in [0, 2]$ with $d(\mathfrak{M}) = 0$ iff $[\omega]$ is dense in $\mathfrak{M}_*$ for every ω. For type I_∞ and type II factors, $d(\mathfrak{M}) = 2$. For type III factors, $d(\mathfrak{M})$ is monotonically decreasing as $\mathrm{III}_0 \to \mathrm{III}_1$ with $d(\mathfrak{M}) = 0$ for type III_1.[9] As a consequence the statespace of a type III_1 algebra is homogenous, with every state appearing equally mixed. This property marks type III_1 factors as maximally noncommutative since no two states can be distinguished by inner-invariant operators.

[8] Connes (1973)

[9] We can write $d(\mathfrak{M})$ as a function of λ:

$$d(\mathfrak{M}) = 2\frac{1 - \lambda^{1/2}}{1 + \lambda^{1/2}}.$$

See Takesaki (2000), for details.

2.1.3 Hilbert Space Geometry

The preceding discussion has been quite abstract. How do the properties of the modular objects (J, Δ) relate back to the original theme of spatial theory? To help see the connection, we begin by discussing a vivid mathematical analogy between the structure of a C^*-algebra and the complex plane. Any element can be uniquely written in "complex form," $A = H+iK$, where $H = 1/2(A+A^*)$ and $K = i/2(A^*-A)$ are self adjoint operators which can be thought of as the "real" and "complex" parts of A. The self-adjoint operators form the "real line." A quick calculation reveals that the involution operation acts as "complex conjugation" on the algebra, $* : H + iK \mapsto H - iK$.[10] This captures the sense in which A and A^* are conjugate

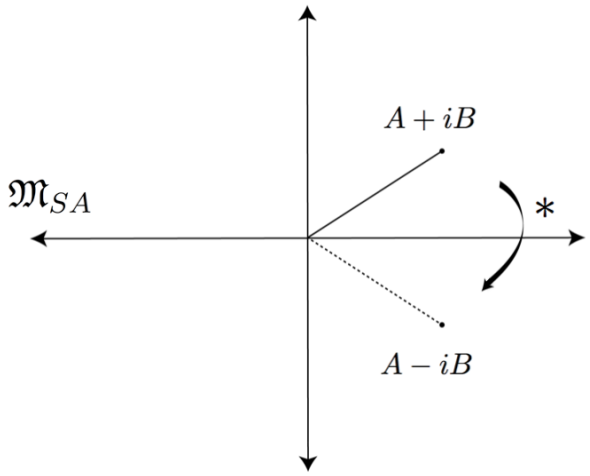


Figure 1: The analogy between $\mathfrak{M}$ and the complex plane

operators and demonstrates that just like the complex numbers, a C^*-algebra is self-similar. Namely, there is a canonical conjugation operation which reflects the algebra across its "real axis." The positive and negative "imaginary halves" of the algebra, while distinct, have identical algebraic properties.

Given a von Neumann algebra with a cyclic, separating vector, we can use tools from modular theory to translate this complex structure to $\mathcal{H}$. Letting $\mathfrak{M}_{SA}$ denote the self-adjoint part of $\mathfrak{M}$, the corresponding subspace $\mathfrak{M}_{SA}\Omega$ can be viewed as a *real* subspace of

[10] Using the definition of $*$ along with the fact that H and K are self adjoint, we have $(H + iK)^* = H^* + (iK)^* = H + \bar{i}K = H - iK$.

$\mathcal{H}$ by taking the closed linear span with respect to the real part of the Hilbert space inner product, $\mathcal{H}^\sharp := \overline{\mathfrak{M}_{SA}\Omega}^{\mathrm{Re}}$. Similarly, $\mathcal{H}^\flat := \overline{\mathfrak{M}'_{SA}\Omega}^{\mathrm{Re}}$ represents the real subspace associated with the self-adjoint part of the commutant. In both cases we find that the complex linear span $\overline{\mathcal{H}^\sharp + i\mathcal{H}^\sharp} = \overline{\mathcal{H}^\flat + i\mathcal{H}^\flat} = \mathcal{H}$.[11] Furthermore, each real subspace has a positive cone $P^\sharp := \overline{\mathfrak{M}_{SA^+}\Omega}^{\mathrm{Re}}$, $P^\flat := \overline{\mathfrak{M}'_{SA^+}\Omega}^{\mathrm{Re}}$, determined by the order structure on each algebra.

Interestingly, the relative positions of the real subspaces $H^\sharp = P^\sharp - P^\sharp$ and $H^\flat = P^\flat - P^\flat$ can be characterized using the modular operators. There is a natural real subspace $\mathcal{H}^\natural = P^\natural - P^\natural$, defined by the self-dual positive cone $P^\natural := \overline{\{AJAJ\Omega \mid A \in \mathfrak{M}\}}^{\mathrm{Re}}$. The cone $P^\natural$ is an invariant of the algebra $\mathfrak{M}$, and characterizes the pair $\mathfrak{M}$, $\mathfrak{M}'$ uniquely (Connes, 1974).[12] It follows from the definitions of the modular objects that $P^\natural = \Delta^{1/4}P^\sharp = \Delta^{-1/4}P^\flat$, thus $\mathcal{H}^\sharp$ and $\mathcal{H}^\flat$ are symmetric with respect to $\mathcal{H}^\natural$. Furthermore, Skau (1980), drawing upon work by Rieffel and van Dale (1977), introduces the operator

$$A_\theta := \frac{|I - \Delta^{1/2}|}{I + \Delta^{1/2}} \tag{2.13}$$

which measures the "tangent" of the angle, θ, between $\mathcal{H}^\sharp$ (alt. $\mathcal{H}^\flat$) and $\mathcal{H}^\natural$. While the positions of $\mathcal{H}^\sharp$ and $\mathcal{H}^\flat$ depend on the choice of vector state Ω, since $P^\natural$ is an invariant of $\mathfrak{M}$, the position of $\mathcal{H}^\natural$ is state independent. Selecting a different generating state adjusts the angle θ.

Drawing upon the analogy with the complex plane, we can think of the polar decomposition, $S = J\Delta^{1/2}$, in terms of "polar coordinates." Recall that any complex number can be expressed as $z = re^{i\theta}$, where $r = |z|$ is a positive number measuring the distance from the origin and $e^{i\theta}$ represents a rotation by θ from the positive real axis. Similarly, the involution operator S can be decomposed into a positive component $\Delta^{1/2} = (S^*S)^{1/2}$ and a rotation J. In this case the rotation is in fact a reflection across $\mathcal{H}^\natural$, taking $\mathcal{H}^\sharp \to \mathcal{H}^\flat$. The modular con-

[11] Note, however, that in general $\mathcal{H}^\sharp$ and $i\mathcal{H}^\sharp$ are not orthogonal, neither are $\mathcal{H}^\flat$ and $i\mathcal{H}^\flat$. Each pair is orthogonal iff the generating state ω is tracial.

[12] In particular, let $\mathfrak{N}$ be some other von Neumann algebra acting on Hilbert space $\mathcal{K}$ with cyclic, separating vector Φ. If $\varphi : \mathfrak{M} \to \mathfrak{N}$ is a $*$-isomorphism, then there exists a unitary operator $U : \mathcal{H} \to \mathcal{K}$ such that, $U(P^\natural_\Omega) = P^\natural_\Phi$ and $\varphi(A) = UAU^*$, for all $A \in \mathfrak{M}$.

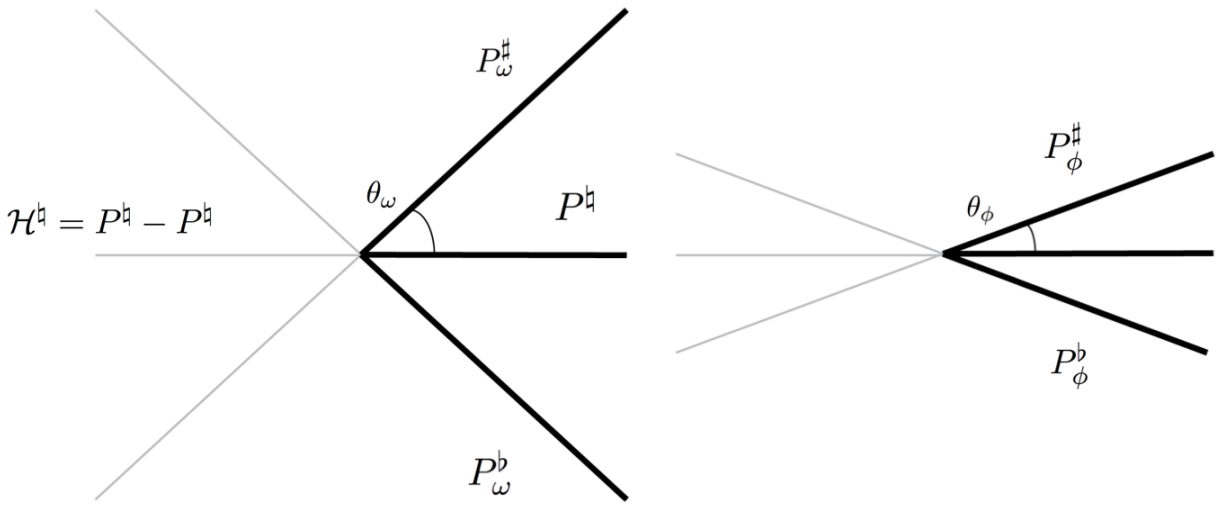


Figure 2: Relative location of real subspaces generated by ω, ϕ.

jugation therefore implements "complex conjugation" on $\mathcal{H}$ with respect to the "real line" $\mathcal{H}^\natural$ associated with the invariant positive cone of $\mathfrak{M}$. It is an anti-isomorphism of the Hilbert space structure. While J reflects vectors, Δ stretches them. The operator $S = J\Delta^{1/2}$ implements complex conjugation with respect to the "real line" $\mathcal{H}^\sharp$, but since S is not antiunitary, it is not an anti-isomorphism of the Hilbert space, but rather an anti-isomorphism with an additional twist corresponding to $\Delta^{1/2}$. This geometric picture vindicates Tomita's original intuition that the algebraic involution does not act isometrically on $\mathcal{H}$.[13]

2.2 The Physics

Modular theory has a wide array of physical applications in QFT and quantum statistical mechanics. Current results indicate that the local algebras in QFT are universally type III_1 factors. As we'll see in Ch. 4.1, modular theory is crucial for proving this result as well as understanding the internal structure of such algebras. Additionally, the KMS property of the modular automorphism group connects modular theory directly to the study of quantum thermodynamics.[14] But this is only the tip of the iceberg,

[13] If ω is tracial, then $S = J$ and the algebraic involution is isometric.
[14] See Bratteli and Robinson (1981) as well as Haag (1996, Ch. 5).

The far-reaching utility of modular theory in QFT is due to a pair of theorems, the Reeh-Schlieder theorem and the Bisognano-Wichmann theorem. Whenever we have a faithful state over a local C^*-algebra, $\mathfrak{A}(\mathcal{O})$, the GNS construction produces a corresponding representation in which the state is represented by a cyclic, separating vector for the local von Neumann algebra $\mathfrak{R}(\mathcal{O}) := \pi(\mathfrak{A}(\mathcal{O}))''$. This provides a minimal arena for modular theory. The Reeh-Schlieder theorem significantly extends this result, ensuring that the vacuum state (or more generally, any state analytic for the energy) will be cyclic and separating for *every* local von Neumann algebra. Thus we can view the entire net of local von Neumann algebras as acting on the same Hilbert space with a common cyclic, separating vector Ω.

Microcausality entails that $\mathfrak{R}(\mathcal{O}') \subset \mathfrak{R}(\mathcal{O})'$, thus the anti-isomorphism implemented by the modular conjugation operator $J\mathfrak{R}(\mathcal{O})J = \mathfrak{R}(\mathcal{O})'$ takes on a geometric character. If duality holds, $\mathfrak{R}(\mathcal{O}') = \mathfrak{R}(\mathcal{O})'$, and the connection is even tighter — the modular conjugation maps $R(\mathcal{O})$ anti-isomorphically onto the algebra of its spacelike complement. So in the spacetime setting of QFT, J starts to look like a reflection map of some kind. Meanwhile, the KMS property shows that Δ^{it} shares many properties with local equilibrium dynamics. The Bisognano-Wichmann theorem makes this pair of analogies mathematically and physically precise within certain physical contexts, establishing a direct link between modular theory and spacetime structure. It is this theorem that stands behind many of the most important applications of modular theory in AQFT.

2.2.1 The Bisognano-Wichmann Theorem

In the case where the net of observable algebras is generated by an underlying set of Wightman fields, Bisognano and Wichmann (1975) prove that the modular objects associated with certain spacetime regions in the vacuum sector can be geometrically interpreted as implementing Poincaré transformations. The regions in question are spacelike wedges — un-

bounded regions, $\mathcal{W}$, formed by a two-dimensional spatial plane and two connected lightlike hypersurfaces.[15]

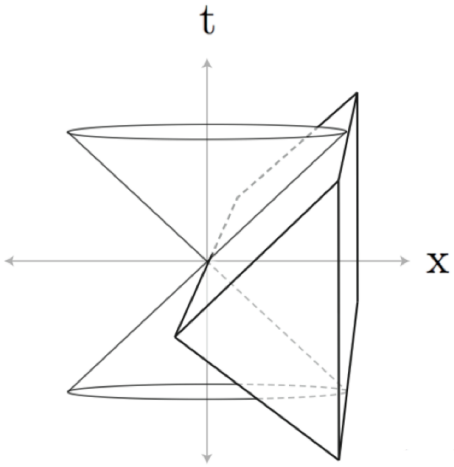


Figure 3: A spacelike wedge

Leting $\Delta_{\mathcal{W}}^{it}$ and $J_{\mathcal{W}}$ denote the modular unitaries and modular conjugation associated with an arbitrary wedge-algebra, $\mathfrak{R}(\mathcal{W})$, in the vacuum state, we have the following:

Theorem 2.3 (Bisognano-Wichmann)**.** *Given a Wightman field theory in the vacuum sector,*

$$\begin{aligned} \Delta_{\mathcal{W}}^{it} &= U(\Lambda_{\mathcal{W}}(2\pi t)),\ \forall t \in \mathbb{R} \\ J_{\mathcal{W}} &= \Theta_1 \end{aligned} \tag{2.14}$$

where $\Lambda_{\mathcal{W}}(2\pi t)$ is the subgroup of Lorentz boosts leaving the wedge region invariant, and Θ_1 is a P_1CT operator, which reverses the direction of time, conjugates charge, and reverses one spatial direction (namely the direction perpendicular to the edge of the wedge).[16]

[15] In standard Minkowski coordinates, the right wedge is defined by $\mathcal{W}_R = \{x \in \mathbb{R}^{1+3} : x^1 > |x^0|\}$. The family of wedge regions consists of all Poincaré transformations of $\mathcal{W}_R$.

[16] See Bisognano and Wichmann (1976) for the proof. The theorem follows from essentially the same assumptions which are used in Jost's constructive proof of the PCT theorem, indicating a tight connection between the two results. Indeed, the PCT theorem for Wightman field theories follows directly from the Bisognano-Wichmann theorem. Since by assumption the theory is covariant with respect to continuous Poincaré transformations, we can easily define a full PCT operator $\Theta := JU(R_\pi)$ where $U(R_\pi)$ is a unitary operator implementing a rotation by π in the plane spanned by the remaining two spatial directions. We'll explore this connection in more detail in the next chapter.

Within the setting of AQFT, we currently do not have a proof of the Bisognano-Wichmann theorem from the Haag-Kastler axioms alone. We do have good reasons to expect that the theorem will hold quite broadly, however. While counterexamples are known to exist, there are none which satisfy both Poincaré covariance and the spectrum condition.[17] Furthermore, Borchers (1992) establishes a tantalizingly close partial result. Assuming only translation covariance and the spectrum condition, he proves that the modular unitaries must act as Lorentz boosts in the characteristic two plane of the wedge — the triangular shaped cross-section orthogonal to the edge of the wedge. For 2-dimensional QFTs and conformal field theories, this result immediately leads to a full version of the Bisognano-Wichmann Theorem.[18]

From a technical standpoint, what is needed is a better understanding of how the modular data associated with a local algebra is related to the modular data of its subalgebras. This has led to the intense study of modular inclusions, situations where the modular operators of a local algebra continue to act as modular operators when restricted to a subalgebra. Both the split and nuclearity conditions (see Ch. 4.1) put constraints on such inclusions, hence many mathematical physicists hope that these assumptions might provided the missing key.[19]

Kuckert (1997) attacks the problem from a slightly different angle. He shows that if one assumes that the wedge modular operators act in any kind of geometric fashion at all, then they must act in the right way, as wedge preserving boosts. The relevant notion of geometric action is extremely weak — the modular conjugation is required to map local algebras to algebras which can be localized in some open region. (It is not necessary that this region have similar geometry or even be simply connected.) How this assumption is related to the

[17] Yngvason (1994) constructs a non-Poincaré covariant net in which the wedge modular unitaries act geometrically in the characteristic two plane, but do not act locally in the direction parallel to the edge of the wedge.

[18] These proofs crucially rely on additional scaling symmetries which are not present in general QFTs. Interestingly, any QFT with a mass gap, including the Standard Model, becomes conformally invariant in the short distance scaling limit. Mund has used this fact to show that any QFT containing only massive particles and satisfying cluster decomposition (or more precisely, which is asymptotically abelian), must obey the Bisognango-Wichman relations. Since massless particles appear to be a necessary component of any interacting gauge theory, this proof still falls short of the grail.

[19] As we'll see in Ch. 3, they also play a small but important role in the PCT and spin-statistics theorems.

split and nuclearity conditions remains an open question. It might also be possible to provide independent justification for Kuckert's symmetry assumption.

In the interim, much progress has been made in AQFT by assuming portions of the Biognano-Wichmann theorem. Three geometric conditions inspired by the theorem which have led to a number of interesting results are *modular covariance*, *geometric modular action*, and *modular stability*. In the following subsection, we will survey a number of their applications.

Modular covariance focuses on the covariance properties associated with the modular automorphism group. In its strongest form, modular covariance requires that for any wedge, $\mathcal{W}$, and any open spacetime region, $\mathcal{O}$, in the vacuum sector,

$$\Delta_{\mathcal{W}}^{it}\mathfrak{R}(\mathcal{O})\Delta_{\mathcal{W}}^{-it} = \mathfrak{R}(\Lambda_{\mathcal{W}}\mathcal{O}), \tag{2.15}$$

i.e. the wedge modular unitaries act as wedge-preserving Lorentz boosts on the entire net of local algebras. This amounts to assuming the first half of the Bisognano-Wichmann theorem. Various weakenings of this condition are surveyed in Davidson (1995).

Geometric modular action (GMA), originally proposed by Buchholz et al. (2000), focuses on the reflection properties of the modular conjugation. The GMA requires that for any physically reasonable state, the collection of wedge algebras must be stable under the adjoint action of the associated modular involutions. Thus for any pair of wedges $\mathcal{W}_1$, $\mathcal{W}_2$, there must be a third wedge $\mathcal{W}_3$, such that

$$J_{\mathcal{W}_1}\mathfrak{R}(\mathcal{W}_2)J_{\mathcal{W}_1} = \mathfrak{R}(\mathcal{W}_3). \tag{2.16}$$

The GMA has the advantage of not making any a priori assumptions about background spacetime isometries, opening up the possibility of extending results to curved spacetimes which have sensible wedge geometry.

The closely related modular stability condition requires the wedge unitaries $\Delta_{\mathcal{W}}^{it}$ to be contained in the group generated by the wedge modular conjugations, $\mathcal{J} =< J_{\mathcal{W}} >$. In Minkowski spacetime, modular stability in conjunction with the GMA condition entails modular covariance (Buchholz et al., 2000, Thm. 5.1.2).

2.2.2 Modular Theory and Spacetime Structure

Within the last two decades there has been an explosion of research activity on QFT-related applications of modular theory in the mathematical physics literature. These include, but are not limited to, quantum thermodynamics, the structure theory of type III von Neumann algebras, entanglement, generalizations of quantum logic and probability theory for field systems, the infrared problem, and no-go theorems for modal interpretations of QFT. It would probably not be an exaggeration to say that virtually every major topic of foundational interest in QFT is linked to modular theory in one way or another. Here we will focus primarily on applications of the Bisognano-Wichmann linking modular theory and spacetime structure.

Unruh effect. In both axiomatic and Lagrangian QFT calculations imply that a uniformly accelerating observer in the Minkowski vacuum will observe thermal black-body radiation with a temperature proportional to her acceleration: $T = a/2\pi k_b$. This result is widely known as the *Unruh effect*.[20]

A wedge region can be interpreted as the region of spacetime causally connected to such a uniformly accelerating observer, the so-called *Rindler Wedge*. If the modular unitaries implement wedge preserving boosts, then the flow of the modular automorphism group will be proportional to the flow of proper time experienced by the accelerating observer. With respect to this flow, the vacuum state is an equilibrium KMS state with a non-zero temperature proportional to the observer's acceleration. Sewell (1980) has proposed this link

[20]Unruh (1976)

as a physical explanation of the Unruh effect.[21] The effect arises because unlike her inertial counterpart, the accelerating observer experiences causal horizons, limiting her access to the full algebra of observables across spacetime. As a result, the restriction of the global vacuum state to the wedge is a thermal KMS state. The non-zero temperature reflects her ignorance of the full global state. If this is deemed a virtuous explanation, then theoretical and experimental support for the Unruh effect can be viewed as support for modular covariance. Some physicists even hope that it might be possible to prove the Bisognano-Wichmann Theorem from the KMS condition directly, although a number of technical obstacles remain.[22]

This thermodynamic link also lies behind the ambitious proposal of Connes and Rovelli (1994) to use modular theory to help solve the problem of time in generally covariant theories. Their *thermal time hypothesis* interprets the Unruh temperature as the ratio between an observer's proper time and the background flow of *thermal time* which is identified with true physical time. We will examine this fascinating idea in more detail in Ch. 5.

PCT and Spin-Statistics Theorems. Modular covariance is a crucial ingredient in extant algebraic proofs of the PCT and spin-statistics theorems.[23] The PCT theorem says that any local, Lorentz invariant QFT will be invariant under a global reflection symmetry that flips spatial handedness (P), conjugates charge (C), and reverses the direction of time (T). Prima facie, it is difficult to understand why this apparently ad hoc combination of operations always yields a perfect symmetry of nature. In the next chapter, I argue that the key to understanding the theorem lies in recognizing that PCT symmetry is not a product of separate symmetry operations at all. Rather it is implemented by a single operation, a total reflection of the theory's statespace. As a result, the symmetry reverses all discrete orientation structures employed by the theory. Crucially, it is the Bisognano-Wichmann theorem that supplies the link between reflections on statespace and reflections on spacetime.

[21] See Earman (2011) for a philosophical critique of this explanation.
[22] See Haag (1996), §V.4.
[23] See Borchers and Yngvason (2000), Guido and Longo (1994), and the references therein.

The closely related spin-statistics theorem shows that in relativistic QFT, particles with integer spin must obey symmetric Bose-Einstein statistics while particles with half-integer spin must obey anti-symmetric Fermi-Dirac statistics. In some versions of the theorem the PCT theorem appears as a crucial lemma, in others it is a direct corollary. This leads to an explanatory chicken-and-egg problem. In the algebraic approach, both theorems are corollaries of the Bisognano-Wichmann theorem, pointing towards a possible resolution — neither theorem is explanatorily primary, they both follow independently from essentially the same assumptions. We leave this as an open question for future investigation.

Spacetime Reconstruction. The algebraic proof of the PCT theorem uses the geometric properties of $J_{\mathcal{W}}$ and $\Delta_{\mathcal{W}}$ to construct a representation of the Poincaré group on $\mathcal{M}$.[24] This raises the question of whether or not it is possible to reconstruct the metrical properties of spacetime from local algebraic data. Summers and White (2003) and Buchholz et al. (2000) provide an affirmative answer for spacetimes with a sufficiently large symmetry group.

Consider a net of von Neumann algebras assigned open regions of a manifold whose metrical and conformal structure is unspecified. Assume the net satisfies isotony and there is a common cyclic, separating vector for all of the algebras. If there is some subnet such that the adjoint action of the associated modular conjugations is a net automorphism (i.e. if an abstract version of the GMA condition holds), Summers et al. show that if the modular conjugations of the subnet exhibit certain additional algebraic relations, it is possible to identify the subnet as a net of wedge algebras on a particular background spacetime.

Let $\mathcal{J}$ denote the group generated by the modular conjugations of the selected subnet. Since the generating set, $\{J_i\}$, satisfies the abstract analogue of the GMA condition, it forms an *invariant* generating set of involutions for $\mathcal{J}$.[25] This is the starting point for *absolute geometry*, which characterizes various metric spaces in terms of their corresponding group of motions. The elements of the generating set correspond geometrically to *lines* in the

[24]The key idea is that any motion in Minkowski space can be generated by reflections about spacelike lines. Once one proves that $J_{\mathcal{W}}$ implements a P_1CT reflection around the edge of the wedge, a representation of the Poincaré group can be directly constructed.

[25]That is, $\mathcal{J}$ is the smallest group containing $\{J_i\}$ and $J\{J_i\}J^{-1} \subseteq \{J_i\}$ for all $J \in \mathcal{J}$.

group plane. Two lines $J_1, J_2 \in \{J_i\}$ are perpendicular iff their product is an involution. *Points* are given by elements that are expressible as the product of two perpendicular lines, i.e. $P := J_1 J_2$ such that $J_1 J_2 \in \mathcal{J}$. Two lines are parallel if they share no common point nor common perpendicular line. Continuing in this fashion, if the generating set satisfies certain axioms, it is possible to identify $\mathcal{J}$ as the group of motions for a particular space.[26] Different axioms give rise to different spacetimes. To date, the method has only been used to construct examples of Minkowksi and de Sitter spacetimes in various dimensions, but in principle it can apply to any spacetime with a large group of isometries.

Model Building. The Bisognano-Wichmann relations also form the basis of a number of new algebraic construction techniques including modular localization (Guido and Longo, 2002), polarization-free generators (Lechner, 2008), and warped convolutions (Buchholz and Summers, 2008). The last method has produced some of the first examples of non-trivial interacting theories in 4d spacetime.[27]

All of these techniques start from the same observation: a QFT is said to be *asymptotically complete* if it possesses a complete particle interpretation.[28] Any such theory determines a pair of unique structures — a representation of the Poincaré group, $U(\mathcal{P})$, satisfying the spectrum condition and an algebra, $\mathfrak{G}$, such that for a fixed wedge, $\mathcal{W}$,

$$\begin{aligned} U(a,\Lambda)\mathfrak{G}U(a,\Lambda)^{-1} &\subset \mathfrak{G}, \text{ if } (a,\Lambda)\mathcal{W} \subset \mathcal{W} \\ U(a,\Lambda)\mathfrak{G}U(a,\Lambda)^{-1} &\subset \mathfrak{G}', \text{ if } (a,\Lambda)\mathcal{W} \subset \mathcal{W}' \end{aligned} \tag{2.17}$$

[26] Summers and White (2003) consider the example of 3d Minkowski spacetime. From the standpoint of projective geometry, 3d Minkowski spacetime can be viewed as an affine space whose plane at infinity is a hyperbolic projective-metric plane. They impose six algebraic constraints on $\{J_i\}$ which are sufficient to ensure that $\mathcal{J}$ is isomorphic to the group of motions of the hyperbolic, projective-metric plane. (Coxeter's theorem shows that every motion of the hyperbolic plane is generated by a product of orthogonal line reflections which correspond to the generating involutions $\{J_i\}$.) This plane is then embedded in a 3d projective space and identified as the plane at infinity, yielding the characterization of Minkowski space. For an introduction to absolute geometry, see Bachmann et al. (1986).

[27] Summers (2012a, §6) contains an excellent survey of these methods.

[28] Formally, this requires that the set of asymptotically free incoming and outgoing states in Haag-Ruelle theory span the entire Hilbert space. Physically, this means that the full energy content of the theory can be recovered from the asymptotic particle configurations.

where (a, Λ) is an arbitrary element of $\mathcal{P}$. Conversely, given a pair $(U(\mathcal{P}), \mathfrak{G})$ satisfying these conditions, we can define a net of local observable algebras by setting $\mathfrak{R}(\mathcal{W}) = U(a,\Lambda)\mathfrak{G}U(a,\Lambda)^{-1}$ and $\mathfrak{R}(\mathcal{O}) = \bigcap \mathfrak{R}(\mathcal{W})$ such that $\mathcal{O} \subset \mathcal{W}$. The resulting net satisfies isotony, covariance, and microcausality. This reduces the construction problem to finding suitable pairs $(U(\mathcal{P}), \mathfrak{G})$, which Summers (2012a) refers to as "germs of quantum field models." Each of the three methods mentioned above uses a different strategy to build germs. Inverting the standard Bisognano-Wichmann derivation, modular localization exploits the PCT operator and appropriate Lorentz boosts from $U(\mathcal{P})$ to construct wedge modular operators which define $\mathfrak{G}$.[29] The second approach, pioneered by Schroer and Weisbrock (2000), uses special wedge-localized elements of the Zamolodchikov-Faddeev algebra to construct low dimensional models with factorizing S-matrices. The method of warped convolutions employs algebraic deformations of free field algebras to generate $\mathfrak{G}$.

QFT in Curved Spacetime. Both the spectrum and covariance axioms make essential use of a background representation of the Poincaré group. In curved spacetime, there is no such representation. To make matters worse, generic solutions of the Einstein field equations have no global isometries. Interestingly, Weisbrock (1992) has proven a converse to Borcher's theorem on the 2d geometric properties of the wedge modular objects. In Minkowski spacetime, any QFT in which $\Delta_{\mathcal{W}}^{it}$ acts as a boost in the characteristic two-plane and $J_{\mathcal{W}}$ acts as a $\mathrm{P_1CT}$ operator satisfies the spectrum condition and has a translation invariant vacuum state. This deep connection has lead to a number of proposals for curved spacetime generalizations of covariance and the spectrum condition. For instance, Summers (2012b) employs modular stability and the GMA condition to provide an intrinsic algebraic characterization of the vacuum state. If this proves to be a fruitful line of research, a future more general version of the Haag-Kastler axioms might incorporate the Bisognano-Wichmann relations or similar geometric constraints on the modular operators as a fundamental assumption.

[29] It is worth noting that the construction relies heavily on the geometric tools from §2.1.3.

Chapter 3

PCT Symmetry and Statespace Geometry

A striking feature of our most basic physical laws is that they appear to be blind to any distinction between past and future. Virtually all of the serious candidates for fundamental laws of nature have been time reversal invariant — for any nomologically possible world, there is an otherwise identical nomologically possible world where the direction of time (and motion) are reversed. If there is even a fact of the matter regarding the direction of time in worlds governed by such laws, it must be grounded in something other than the laws themselves. Surprisingly, this is not generally true for quantum field theories. It is possible to write down physically reasonable QFTs which are not time reversal invariant, and as James Cronin and Val Fitch demonstrated in 1964, weak nuclear interactions in the actual world are described by such a theory.[1]

Unfortunately for philosophers of time, this kind of asymmetry cannot be leveraged into an explanation of time's arrow. While QFTs may fail to be symmetric under simple time reversal, a deep foundational theorem ensures that there is always a more complicated time reversal symmetry present. The *PCT Theorem* states that any local, relativistic QFT must

[1] In work that would win them the 1980 Nobel Prize, Fitch and Cronin observed that neutral kaons transform into their antiparticle partners at a different rate than the reverse process.

be invariant under a combined symmetry operation that reverses time, flips parity (spatial handedness), and conjugates charges (e.g. electric charge, weak hypercharge, color charge). Since particles and antiparticles carry opposite charge, the net effect of charge conjugation is to swap matter and antimatter. In a PCT invariant theory, every nomologically possible world has a dopplegänger where the future is the past, right is left, and you and I are made out of antiparticles.[2]

It is far from obvious why this should be the case. Why does this seemingly ad hoc assemblage of operations always yield a perfect symmetry? The theorem becomes even more puzzling when one recasts it in terms of the effects of a PCT transformation on possible global orientation structures a QFT might make use of — temporal orientation, spatial orientation, and charge orientation. If a theory fails to be time reversal invariant, say, it does so because its laws somehow pick out a distinguished temporal direction. The PCT theorem then (roughly) says that no well-behaved QFT can make use of one kind of orientation independently of the other two kinds. A theory can single out a preferred temporal orientation, but only up to a choice of charge sign and spatial handedness. This is rather mysterious, for as Hilary Greaves puts it, these orientation structures appear to be "paradigm cases of distinct existences."[3] Indeed, in relativistic spacetime one can show that spatial orientation and temporal orientation are mathematically independent; a choice of one does not fix the other. While their relationship to charge orientation is less clear cut, charge superselection structure arises from internal gauge symmetries associated with particular forces the charges couple to, and these internal symmetries have no direct connection to spacetime structure.[4]

Such a connection appears essential, however, in order to unify quantum mechanics and special relativity. Straightforward extensions of Schrödinger wave mechanics to relativistic

[2]As in the case of T invariance, PCT invariance is often interpreted as indicating that these apparently distinct possibilities are in fact different representations of the same physically possible world.

[3]Greaves (2010), p. 11.

[4]Here the Coleman-Mandula theorem (Coleman and Mandula, 1967) gives us prima facie support for a kind of independence result. The theorem says that in the absence of supersymmetry, there is no non-trivial way to combine internal gauge symmetries and external spacetime symmetries. This result makes a necessary connection between C and PT symmetries all the more puzzling.

spacetime generate non-physical negative energy states. A series of no-go results show that it is impossible to excise these states while maintaining locality and Lorentz invariance. Interestingly, in quantum field theory it becomes possible to reinterpret negative energy states as positive energy antiparticle states. This trick only works to restore Lorentz invariance, though, if there is an exact correspondence between particles and antiparticles — they must be indistinguishable except for their charge. The PCT theorem accounts for this, explaining why particle/antiparticle pairs have the same mass, spin, and lifetime. It is therefore a result with profound foundational significance. A breakdown of PCT symmetry would signal the need for a radical overhaul of our understanding of relativistic quantum theory.

Despite its importance, the physical basis for the PCT Theorem remains obscure. Part of the problem is that there are several different versions of the theorem with different starting assumptions. Many of the more technical assumptions don't have a clear physical interpretation, making comparisons between the various proofs challenging. To compound this difficulty, the theorems are couched within the different mathematical frameworks of Lagrangian, constructive, and algebraic QFT. The Lagrangian proofs, while more physically transparent, lack mathematical rigor, whereas the rigorous axiomatic proofs are more physically opaque. This state of affairs has prompted recent work by Hilary Greaves and Teruji Thomas on a rigorous Lagrangian theorem based on ideas from John Bell's 1955 proof.[5] In this chapter I attack the problem from the opposite direction, by looking for a more physically perspicuous interpretation of the axiomatic proofs.

As we'll see, tools from modular theory can help shed light on how local algebraic symmetries can give rise to global geometric ones. The main insight gained is the following: in a local, Lorentz invariant QFT temporal, spatial, and charge orientation structures are less independent than they first appear. This can already be seen in the case of relativistic spacetime where the presence of an additional structure, a *total orientation*, renders temporal and spatial orientations mutually interdefinable. If a theory's laws somehow picked out

[5]Bell (1955), Greaves and Thomas (2012)

a preferred total orientation, then handedness and the direction of time would be bound at the hip. While this link is insufficient to account for the full pattern of connections implied by the PCT theorem, it has a more abstract cousin which can.

As it turns out, each local algebra, $\mathfrak{R}(\mathcal{O})$, is an intrinsically oriented object. The choice of an orientation amounts to the specification of a Lie bracket compatible with the underlying Jordan product, giving rise to a geometric orientation on the corresponding GNS Hilbert space. This orientation structure is critical for encoding the manner in which quantum observables generate local symmetries. Since they are self-adjoint, the local physical quantities are always invariant under a general reflection symmetry which reverses the local statespace orientation.The Haag-Kastler axioms, in turn, establish a tight correspondence between local and global statespace structure, while the Bisognano-Wichmann theorem links the geometry of statespace to the geometry of spacetime. These connections make it possible to piece together local reflections into a global PCT operator. Thus, much in the same way that a total spacetime orientation constrains other spatiotemporal structures, a statespace orientation establishes necessary connections between temporal, spatial, and charge orientations. Rather than the product of three separate operations, PCT symmetry is revealed to be the physical manifestation of a single global reflection of the theory's statespace.

The plan of the chapter is as follows: I begin in §3.1 with a warmup example designed to clarify the relationship between various orientation structures on Minkowski spacetime. This serves both to give a flavor for the kind of explanation I'm looking for as well as to establish several technical facts which will prove important down the line. In §3.2 I discuss how AQFT explains charge structure, leading to a statement of the PCT theorem in §3.3 The main argument is contained in §3.4, where I outline a full proof of the PCT theorem, focusing on hidden connections between orientation structures on statespace and spacetime. In §3.5 I compare this account to a different explanation of the PCT theorem offered by Greaves (2010). The two accounts diverge significantly. Whereas Greaves maintains that

the theorem is essentially a relativistic result, in §3.6 I argue that the theorem crucially relies on central assumptions from both relativity and QM.[6]

3.1 Spacetime Orientation

Since two out of three pieces of the PCT puzzle are spatiotemporal in origin, as a warmup, we'll ignore the complications associated with charge conjugation and focus on the relationship between symmetries and orientation structure in Minkowski spacetime.

The group of continuous, metric-preserving symmetries of Minkowski spacetime is the Poincaré group, $\mathcal{P}$. Its elements can be classified as either *connected* or *disconnected*. Connected symmetries include temporal and spatial translations (constant shifts in a timelike or spacelike direction), spatial rotations (rotations in a spacelike hyperplane), and Lorentz boosts (generalized spatiotemporal rotations which are interpreted as transitions to a constantly accelerating frame of reference). Such symmetries can be infinitesimally generated as the flow along integral curves of some vector field on $\mathcal{M}$. As a result, they are topologically connected to the identity transformation. Disconnected symmetries cannot be generated in this fashion. Such symmetries correspond to spatiotemporal reflections and include time reversal, T, spatial inversion, P, and their product, PT.[7]

The disconnected symmetries can be classified by their effect on various orientation structures that can be added to $(\mathcal{M}, \eta)$. A *temporal orientation*, τ, is a continuous labeling of

[6]Material from this chapter was originally presented at the *Quantum Time* workshop at the University of Pittsburgh in April 2014.

[7]$\mathcal{P}$ is a real, ten dimensional noncompact Lie group equal to the semi-direct product of the translations, $\mathcal{T} = \mathbb{R}^{1,3}$, and the Lorentz group, $\mathcal{L} = O(1,3)$,

$$\mathcal{P} = \mathcal{T} \rtimes \mathcal{L}.$$

Thus $\mathcal{P}$ is isomorphic to $\mathbb{R}^{1,3} \times O(1,3)$ with an additional group action of the normal subgroup of generalized rotations on the translations. A generic element can be represented by $\{a, \Lambda\}$, where a is a translation vector and Λ is an element of the Lorentz group. The group multiplication law is given by

$$\{a_1, \Lambda_1\}\{a_2, \Lambda_2\} = \{a_1 + \Lambda_1 a_2, \Lambda_1 \Lambda_2\}.$$

future and past lobes of the lightcones at every point in spacetime. Mathematically, this can be represented by an equivalence class of co-aligned continuous timelike vector fields $[t^a]$.[8] In Minkowski spacetime, there are only two possible temporal orientations, τ and its opposite $\overline{\tau}$. A time reversal transformation, T, is an isometry from $(\mathcal{M}, \eta, \tau)$ onto $(\mathcal{M}, \eta, \overline{\tau})$.[9]

A *spatial orientation*, ϵ, is a continuous labeling of right- and left-handed triads of linearly independent spacelike 4-vectors. This can be represented by choosing a nowhere vanishing 3-form (a completely antisymmetric (0,3)-rank tensor field) that is orthogonal to some timelike vector field, t^a. This defines a consistent notion of spatial handedness across spacetime.[10] Any two such 3-forms related by a strictly positive function will pick out the same orientation, thus we define ϵ as an equivalence class of 3-forms related by strictly positive functions. Like a temporal orientation, there are two possible spatial orientations on Minkowski spacetime. A parity transformation, P, is an isometry from $(\mathcal{M}, \eta, \epsilon)$ onto $(\mathcal{M}, \eta, \overline{\epsilon})$ that reverses all three independent spatial directions.[11]

While both of these orientation structures require the metric, there is a third kind of orientation which depends solely on the underlying manifold structure. A *total orientation*, ε, is a continuous labeling of right- and left-handed tetrads of linearly independent 4-vectors, represented by an equivalence class of nowhere vanishing 4-forms related by strictly positive functions.[12] This defines a consistent notion of spatiotemporal handedness across spacetime. While individually P and T transformations reverse total orientation, together PT preserves

[8] We say that two timelike vectors x^a, y^a are *co-aligned* if $x^a y_a > 0$. This defines an equivalence relation on the set of timelike vectors at every point in $\mathcal{M}$. (See Malament (2012), Prop. 2.2.2.) Two timelike vector fields are co-aligned if they are co-aligned at every point.

[9] In $O(1,3)$, the time reversal operator is given by $T = \mathrm{diag}(-1,1,1,1)$.

[10] Consider an appropriate 3-form field e and a triad of orthogonal, spacelike basis vectors (b^1, b^2, b^3) at a point $p \in M$. The scalar value $e(b^1, b^2, b^3) \neq 0$ iff e is nonzero at p. Thus e splits the set of all spacelike basis vectors at p into two classes, those for which $e(b^1, b^2, b^3)$ is positive and those for which it is negative. These classes can be used to represent right- and left-handedness.

[11] In $O(1,3)$, the parity operator is given by $P = \mathrm{diag}(1,-1,-1,-1)$, the PT operator by $\mathrm{diag}(-1,-1,-1,-1)$. Note that in spacetimes with odd dimension, this definition doesn't make sense since P and PT must have determinant -1 and 1 respectively. In this case, we define a parity transformation as ϵ-reversing isometry which reflects $n-1$ spatial directions.

[12] This works in exactly the same way as spatial orientation, except now we need a 4-form since we're assigning spatiotemporal handedness tetrads of 4-vectors.

ε, while reversing τ and ϵ.[13] To summarize, we have:

$$\begin{aligned} T &: (\mathcal{M}, \eta, \tau, \epsilon, \varepsilon) \to (\mathcal{M}, \eta, \overline{\tau}, \epsilon, \overline{\varepsilon}) \\ P &: (\mathcal{M}, \eta, \tau, \epsilon, \varepsilon) \to (\mathcal{M}, \eta, \tau, \overline{\epsilon}, \overline{\varepsilon}) \\ PT &: (\mathcal{M}, \eta, \tau, \epsilon, \varepsilon) \to (\mathcal{M}, \eta, \overline{\tau}, \overline{\epsilon}, \varepsilon) \end{aligned} \tag{3.1}$$

Let us suppose for a moment that we have a PT theorem, a result showing that any well-behaved QFT must be invariant under PT transformations. Alternatively we can view this as a statement that no well-behaved QFT can make use of a temporal or spatial orientation independently of the other. What could possibly account for this?

One option is that in Minkowski spacetime, temporal direction and spatial handedness might be definitionally linked. Then if we had some general argument for dynamical reversability (perhaps drawing on the fact that time translations are implemented by unitary transformations in QFT), we could leverage it into an argument for PT invariance. Unfortunately, this is not the case. The two concepts are mathematically independent, selecting one does not fix the other. This can be formalized as follows:

Fact 3.1. *The structure* $(\mathcal{M}, \eta, \tau)$ *does not naturally define* ϵ*. Nor does* $(\mathcal{M}, \eta, \epsilon)$ *naturally define* τ*.*[14]

Here we have appealed to the concept of *natural definability*. An important mathematical tool that we will employ throughout this paper, natural definability captures the sense in which one mathematical structure can be defined in terms of another without making any arbitrary choices.[15] There is an important general connection between natural definability and symmetries. If one structure $(\mathcal{M}, X)$ naturally defines another $(\mathcal{M}, Y)$ over the same

[13] As a topological group, $\mathcal{P}$ has four connected components: $\mathcal{P}^{\uparrow}_{+}$, $\mathcal{P}^{\uparrow}_{-}$, $\mathcal{P}^{\downarrow}_{-}$, $\mathcal{P}^{\downarrow}_{+}$. The labels indicate whether elements preserve $(\uparrow, +)$ or reverse $(\downarrow, -)$ the temporal and total orientation respectively.

[14] *Proof.* $\mathrm{Aut}(\mathcal{M}, \eta, \tau) = \mathcal{P}^{\uparrow} = \mathcal{P}^{\uparrow}_{+} \cap \mathcal{P}^{\uparrow}_{-}$, while $\mathrm{Aut}(\mathcal{M}, \eta, \epsilon) = \mathcal{P}^{\uparrow}_{+} \cap \mathcal{P}^{\downarrow}_{-}$. Specifically P reverses ϵ while preserving τ, while P reverses τ while preserving ϵ. So $\mathrm{Aut}(\mathcal{M}, \eta, \tau) \nsubseteq \mathrm{Aut}(\mathcal{M}, \eta, \epsilon)$ and $\mathrm{Aut}(\mathcal{M}, \eta, \epsilon) \nsubseteq \mathrm{Aut}(\mathcal{M}, \eta, \tau)$.

[15] This is the cousin of *explicit definability* from model theory. In category theory, the notion of a *natural transformation* can also be used to help precisify the concept.

base set, $\mathcal{M}$, then the symmetries of the former are a subset of the latter (i.e. $\mathrm{Aut}(\mathcal{M}, X) \subseteq \mathrm{Aut}(\mathcal{M}, Y)$).[16] This captures a sense in which the structure Y supervenes on the structure $(\mathcal{M}, X)$ — Y is completely and uniquely determined by $(\mathcal{M}, X)$, there can be no difference in Y without a corresponding difference in $(\mathcal{M}, X)$. We will assume that in well-formulated theories, this mathematical supervenience relation mirrors physical supervenience. Thus if a theory's laws (understood as mathematical structures $\{L\}$ on $\mathcal{M}$) naturally define some additional structure Y, then the theory is committed to whatever physical structure Y represents, regardless of whether or not Y features explicitly in the formulation of the laws. We propose this as a technical precisification of the notion of "essential use" — a theory's laws make essential use of a certain structure Y iff $(\mathcal{M}, \eta, \{L\})$ naturally defines Y.[17]

Interestingly, while τ and ϵ are not interdefinable in this sense, once we add a total orientation, they become definitionally linked. In fact, in Minkowski spacetime, any two orientation structures naturally define the third:

Fact 3.2. *The structure $(\mathcal{M}, \eta, x, y)$ naturally defines z, where x, y, z, all distinct, range over possible orientation structures $\tau, \epsilon, \varepsilon$.*[18]

Thus, while τ, ϵ, and ε are largely independent of one another, they are not completely independent. Fixing a temporal orientation by itself does not constrain the choice of a spatial orientation, but if a theory's laws also make essential use of a total orientation, then together the temporal and total orientations naturally define a consistent spatial orientation. It comes along for free. Moreover, reversing the temporal orientation entails a parity inversion,

[16] In the special case of complete, first-order theories, the converse holds as a consequence of the Beth-Svenonius Theorem, making the connection between symmetries and definability even tighter.

[17] See Malament (1977), Swanson and Halvorson (2012), and Clifton (2000) for applications of this principle in the domains of relativity, classical mechanics, and quantum field theory respectively.

[18] *Proof Sketch.* Consider $t^a \in \tau$, $e_{abc} \in \epsilon$, and $e_{abcd} \in \varepsilon$. First take t^a and e_{abc}. The latter must be orthogonal to some timelike vector field s^a co-aligned with t^a. Define the 4-form $e_{abcd} := -s_{[a}e_{bcd]}$. Now take t^a and e_{abcd}. Using the natural contraction operation, define the 3-form $e_{bcd} := t^a e_{abcd}$. Since t^a is timelike, e_{abc} must be orthogonal to some timelike vector field. Finally, take e_{abc} and e_{abcd}. Define the vector field $t^a := e^{abcd}e_{bcd}$. Since e_{abc} is orthogonal to some timelike vector field, t^a must be timelike. In all three cases since the input objects are nowhere vanishing, the defined object is similarly nowhere vanishing. Furthermore, one can show that given different choices of input objects from the same equivalence class, the defined object will be unique up to a positive function, thus the definitions are natural. □

since the opposite temporal orientation naturally defines the opposite spatial orientation (in conjunction with the fixed total orientation). In such a theory, temporal direction and handedness are bound together.

A general argument for dynamical reversibility could then yield a PT theorem. The hitch is that such a theorem bears little resemblance to the PCT theorem that we actually have, not just because of the possibility of PT violation compatible with PCT invariance, but because it rules out the possibility of a theory which is invariant under P or T transformations independently. If $(\mathcal{M}, \eta, \{L\})$ naturally define a total spacetime orientation, then $\mathrm{Aut}(\mathcal{M}, \eta, \{L\}) \subseteq \mathrm{Aut}(\mathcal{M}, \eta, \varepsilon)$ which automatically rules out ε-reversing symmetries like T and P.[19] But this is a distinct theoretical possibility within the scope of the PCT theorem. In fact both quantum electrodynamics and quantum chromodynamics are invariant under P, T, C, PT, CP, and TC individually.

While this kind of explanation fails to give us a PT analogue of the PCT theorem, it was a useful exercise. It illustrates that orientation structures in Minkowski spacetime are not as independent as they initially seem. Moreover, we'll find that a similar explanatory pattern arises in a slightly different setting. In order to represent charges, we need to introduce more abstract geometric tools defined on a QFT's statespace. In §3.4.1 we'll discover that quantum mechanical statespace carries a natural orientation structure defined by the Lie product. This statespace orientation in turn constrains temporal, spatial, and charge orientation structure, providing the key to unlocking the PCT theorem.

[19]This assumes that we identify symmetries of $(\mathcal{M}, \eta, \{L\})$ with automorphisms of the structure. In quantum theories, we generally want to allow for symmetries which are anti-automorphisms of $(\mathcal{M}, \eta, \{L\})$ as well. (These are implemented by anti-unitary mappings.) This opens up a possible loophole if it turns out that P and T transformations are always anti-automorphisms. This is in fact true for T but not for P. Thus the possibility of unitary parity inversion sinks our attempt at recovering a PT analogue of the PCT theorem via a natural total spacetime orientation.

3.2 Superselection Rules and Antimatter

Before we can state the PCT theorem, we need to put charge back into the picture. One of the main achievements of AQFT is a comprehensive picture of superselection rules and corresponding charge structure. This yields an elegant characterization of antimatter, a virtue which makes AQFT uniquely situated to shed light on the mysterious connection between C and PT transformations implied by the PCT theorem.

Charges correspond to conserved quantities associated with particular force laws. Electric charge is the conserved quantity that couples to the electro-magnetic force. Color charge is the conserved quantity that couples to the strong force, etc. Besides satisfying global conservation laws, one of their peculiar features is that in quantum theory they are superselected quantities. A *superselection rule* forbids states which are superpositions of different charges.[20] Such rules have an interesting effect on the theory's statespace. Rather than a single, connected Hilbert space, the physical states of the theory can only be represented using multiple, disconnected Hilbert spaces. These *sectors* represent states with the same global charge, Q, and correspond to unitarily inequivalent representations of the quasilocal algebra of observables, $\mathfrak{A}$. Intuitively we can picture the sectors as a collection of "dynamical islands." Since the sectors are unitarily inequivalent and the dynamics (on the standard view) are assumed to be unitary, the flow generated by the dynamical laws must stay within each island. This reflects the fact that globally, charges are conserved.

Charges have two other characteristic features — they can be *localized* in some subregion of spacetime and every charge appearing in the theory has a well-defined opposite or *conjugate* charge. Particles carrying conjugate charges can annihilate each other, producing pure energy. Conversely, particle/antiparticle pairs can spontaneously spring from the vacuum

[20]Note that in the case of quantum chromodynamics the superselected charges are the total color hypercharge magnitude (Y_c^2) and the total color isospin magnitude (I_c^2). These are gauge invariant Casimir operators associated with the local gauge group $SU(3)$. The standard color charge labels *red*, *green*, *blue*, correspond to non-gauge invariant combinations of the Y_c and I_c components and thus (on the usual interpretation) have no direct physical significance.

state. In AQFT, these features are captured using special mappings of the global observable algebra into itself called *localized transportable endomorphisms*.[21]

Mathematically, the collection of localized transportable endomorphisms has a rich structure, that of a symmetric tensor $*$-category. In particular, the category has a natural tensor product which allows us to define notions of charge composition and conjugate charges. The endomorphisms induce a mapping on global states over $\mathfrak{A}$. Letting ω represent the unique vacuum state, and ϱ a particular localized transportable endomorphism, we can think of $\omega \circ \varrho$ as describing a state with charge Q localized in region $\mathcal{O}$. Its conjugate is defined as the unique endomorphism, $\bar{\varrho}$, such that $\omega \circ \varrho \circ \bar{\varrho}$ is a mixed state containing a component in the vacuum sector. This captures the necessary condition for pair creation/annihilation.[22]

In their groundbreaking analysis of superselection structure, Doplicher, Haag, and Roberts considered charges described by endomorphisms localized in compact spatiotemporal regions.[23] Such charges couple to forces whose strength falls off sharply as a function of distance, like the strong force. They showed that the relevant category of localized transportable endomorphisms is in fact equivalent to the category of representations of $\mathfrak{A}$ which differ from the vacuum within a compact region and are unitarily equivalent to the vacuum everywhere else. The entire charge structure can thus be directly encoded in the locality properties of different inequivalent representations of $\mathfrak{A}$. Different representations correspond to different sectors. Conjugate sectors are defined by the condition that $\pi(\mathfrak{A}) \otimes \bar{\pi}(\mathfrak{A})$ contains a copy of the vacuum representation $\pi_0(\mathfrak{A})$.[24]

[21]Formally, such a morphism is an injective $*$-homomorphism, $\varrho : \mathfrak{A} \to \mathfrak{A}$. It must be localized in some spatiotemporal region $\mathcal{O}$ (i.e. it is trivial in the spacelike complement $\mathcal{O}'$), and it must be possible to shift ρ over to any other similarly shaped region in spacetime using unitary mappings (i.e. for any similar region $\mathcal{O}_2$ there is a localized endomorphism ϱ_2 and a unitary operator U such that $U\varrho(A) = \varrho_2(A)U$, for all $A \in \mathfrak{A}$).

[22]A standard annihilation event is $e^+ + e^- \to \gamma + \gamma$. Since charge is globally conserved and photons are chargeless, any interaction of this kind requires that particles and antiparticles have conjugate charge. If $\omega \circ \varrho \circ \bar{\varrho}$ is a mixed state with a component in the vacuum sector, then the probability of a creation/annihilation event is nonzero according to the Born rule.

[23]For a detailed treatment of DHR superselection theory, see Halvorson and Müger (2006), §8-11 and the references therein.

[24]Their analysis has subsequently been extended by Buchholz and Fredenhagen (1982) to include charges localized in spacelike cones, which are sufficient to describe any Lorentz invariant theory with massive, charged particles. Because of the infrared problem in quantum electrodynamics, we currently do not fully

According to this picture, matter and antimatter states are represented by states in conjugate superselection sectors. This gives rigorous mathematical content to the idea that such states have opposite charge quantum numbers. As argued by Baker and Halvorson (2010), it has the additional advantage of applying to QFTs which lack a clear particle interpretation. In contrast, the standard characterization of antimatter from Lagrangian QFT crucially relies on the emergent particle picture. This suggests that the DHR picture provides a more fundamental description of antimatter.[25]

3.3 The PCT Theorem

With these tools in place, we can turn our attention to the PCT theorem. The theorem can be viewed as a kind of existence proof — given a generic quantum field theory, there always exists an operator implementing a certain global reflection symmetry. But what, exactly, is a symmetry of the theory?

A natural candidate would be automorphisms of $\mathfrak{A}$, but these can wreak havoc on the local algebras, so it is not true in general that any global automorphism is a symmetry. What we need is a notion of a *net automorphism*. One possibility is an automorphism of the global algebra that leaves each local algebra invariant. But this is too restrictive since it rules out the possibility of symmetries which map different local algebras onto one another. These include continuous spacetime symmetries as well as PCT symmetry.

Halvorson and Müger (2006) suggest a solution to this problem. Our net is a partially ordered set ordered by spatiotemporal inclusion (via the isotony axiom). We can let a symmetry of the net be given by a pair (F, α) where F is an order preserving bijection and α is a morphism from $\mathfrak{A}$ to $\mathfrak{A} \circ F$. Roberts and Roepstorff (1969) show that such mappings must preserve transition probabilities between states, hence they must be implementable by

understand the localization properties of massless charged particles. See Buchholz et al. (2013) for recent work on this issue.

[25] See Wallace (2009) for the opposing viewpoint.

either unitary or antiunitary operators.[26] (This result can be seen as an algebraic analogue of Wigner's theorem.)

This leads to a statement of our main result:

Theorem 3.3 (PCT Theorem)**.** *Given a generic quantum field theory (i.e. a net of von Neumann algebras over Minkowski spacetime satisfying the Haag-Kastler axioms) there exists an antiunitary PCT operator, Θ, implementing an automorphism of the net.*

In order for a candidate PCT operator to deserve the title, we require that it satisfy three general constraints:

1. $\Theta\mathfrak{R}(\mathcal{O})\Theta = \mathfrak{R}(-\mathcal{O})$, for every local von Neumann algebra $\mathfrak{R}(\mathcal{O})$,

2. $\Theta U(a,\Lambda)\Theta = U(-a,\Lambda)$, where $U(a,\Lambda)$ is a unitary operator representing a generic connected Poincaré transformation (a represents a translation, Λ represents a boost/rotation),

3. $\Theta\varrho\Theta = \bar{\varrho}$, for any localized transportable charge endomorphism ϱ.

The first constraint tells us that Θ maps each local algebra onto the algebra associated with a region obtained from $\mathcal{O}$ by a full spatiotemporal reflection (in both the space and time coordinates). The second represents the requirement that Θ interact with other spacetime symmetries in the appropriate fashion, namely, by reflecting the translation component of a generic Poincaré transformation. The final constraint reflects the fact that Θ must implement a proper charge conjugation symmetry, mapping each charge sector to its conjugate sector.

Alas, no one knows how to prove the PCT theorem in this form. All known proofs require additional assumptions that go beyond the Haag-Kastler axioms, so at this point what we've stated is really a PCT conjecture. The standard axiomatic proof that is widely cited in the

[26] An operator, V, is *antiunitary* iff it is antilinear and preserves transition probabilities, i.e. $(VA, V, B) = \overline{(A, B)}$. Note that defined in this way, a symmetry of the net need not give rise to an (anti)unitary equivalence between the associated representations. A symmetry only has to intertwine the global and local algebras, it doesn't have to do so pointwise.

literature is originally due to Jost.[27] This version of the theorem only applies in situations where the net of local observable algebras is generated by an underlying system of *Wightman fields*. These fields are structurally similar to the unbounded local field operators appearing in Lagrangian QFT. On the standard interpretation, they do not directly represent physical quantities at all. This is the role of the observables, which correspond to gauge-invariant combinations of field operators.

Jost's proof is a mathematical tour de force, representing a significant conceptual advancement over earlier Lagrangian proofs.[28] Nonetheless, it has several significant drawbacks. First, it functions by placing restrictions on the field operators rather than the observables. Since these do not directly represent quantities, it makes the physical interpretation and justification of these constraints difficult. Second, the proof techniques crucially rely on the fact that the field operators are localized in regions of compact support. It is known that realistic interacting theories can only be modeled using non-local field operators, making it difficult to see how the proof could possibly be generalized to cover these critical cases.[29] Third, it is extremely difficult to diagnose the physical source of PCT invariance from Jost's proof as it hinges on the procedure of analytically extending vacuum correlation functions, a process which lacks clear physical content.

[27] See Jost (1965), also Streater and Wightman (1989) for a comprehensive analysis of the standard axiomatic proof. Bain (2013) contains an excellent philosophical analysis of the structural differences between Jost's proof and its algebraic and Lagrangian counterparts.

[28] Jost's proof relies crucially on the analytic properties of Wightman functions which encode correlations between the Wightman fields in the vacuum state, $\langle\Omega|\phi_1(x_1)\phi_2(x_2)\dots\phi_n(x_n)|\Omega\rangle$. Here's the standard gloss: by a Lemma due to Bargmann, Hall, and Wightman, the spectrum, covariance, and finiteness conditions imply that the Wightman functions have a unique analytic extension to correlation functions of complex variables. These functions transform irreducibly under a representation of the complex Poincaré group $\mathcal{P}(\mathbb{C})_+$ (or its covering group $\tilde{\mathcal{P}}(\mathbb{C})_+$). Unlike the real Poincaré group, in the complex Poincaré group, the identity transformation I is topologically connected to the element $-I$, which implements a total PT reflection. The spectrum condition requires that this reflection be implemented by an anti-unitary operator which conjugates the fields and reverses their order, $\langle\Omega|\phi_1(x_1)\phi_2(x_2)\dots\phi_n(x_n)|\Omega\rangle \mapsto \langle\Omega|\phi_n(-x_n)^*\phi_{n-1}(-x_{n-1})^*\dots\phi_1(-x_1)^*|\Omega\rangle$. The existence of such an operator requires the field algebra to be closed under involution. Additionally, one can show that the complex Wightman functions are analytic on a set of real points, called *Jost points*, which are spacelike separated. Because of this, the Field-locality assumption implies that one can restore the original order of products and so the real Wightman functions are also invariant under a (real) PCT transformation. Since the entire theory can be reconstructed from the n-point Wightman functions, the full theory is also PCT invariant.

[29] See Strocchi (1993), Ch. VI.

In contrast, the axioms of AQFT have a more transparent physical interpretation. Moreover the locality assumptions they make concern observables not fields, providing a possible route to circumvent the no-go results mentioned above. A rigorous, fully general algebraic proof of the PCT theorem would be a boon. While at present no such proof exists, there are also no known counterexamples, and we have rigorous algebraic proofs for QFTs in more limited domains: 2-dimensional Minkowski spacetime (Borchers, 1992), conformal field theories (Brunetti et al., 1993), and asymptotically complete theories of massive particles (Mund, 2001). Interestingly, all of these proofs rely on generalizations of the Bisognano-Wichmann theorem (see Ch. 2.2.1), linking the PCT theorem directly to modular theory. It turns out that one can construct a PCT operator on the full net of observables using wedge modular conjugations. In the next section I will argue that this fact can be used to shed light on the geometric interpretation of PCT symmetry more broadly.

My argument takes the form of an outline of a full proof of the PCT theorem employing the Haag-Kastler axioms and two additional assumptions: modular covariance (see Ch. 2.2.1) and the split condition (see Ch. 4.2.1).[30] The emphasis will be on the geometric interpretation of the involution mapping as a reflection of local statespace and its connection to global spacetime symmetries supplied by the Haag-Kastler axioms and modular covariance. In addition to modular covariance, the heavy lifting is done by three axioms: Poincaré covariance, microcausality, and the spectrum condition. (The split condition plays a supporting role.) It is interesting to note that it was the problem of combining these three constraints that led to the development of quantum field theory in the first place. This suggests that the PCT theorem is closely tied to the unification of quantum mechanics and relativity. As we'll see in §3.6, this is in fact the case.

[30]These are the same assumptions that Brunetti et al. (1993) make in their proof, although the details diverge.

3.4 Dissecting the Theorem

The PCT theorem tells us that every generic (local, relativistic) QFT will be invariant under PCT transformations. The Haag-Kastler axioms give us a good initial characterization of such theories which may have to be supplemented. The lone axiom that explicitly deals with symmetry is the covariance axiom, but this only ensures that the theory is invariant under connected Poincaré symmetries. How can a disconnected symmetry like PCT get into the mix? The answer lies in the structure of the local algebras themselves. Recall that each local algebra is closed under a canonical involution mapping $* : \mathfrak{R}(\mathcal{O}) \to \mathfrak{R}(\mathcal{O})$, such that:

$$\begin{aligned} (A^*)^* &= A \qquad (A+B)^* = A^* + B^* \\ (cA)^* &= \bar{c}A^* \qquad (AB)^* = B^*A^* \end{aligned} \tag{3.2}$$

for any $A, B \in \mathfrak{R}(\mathcal{O})$, $c \in \mathbb{C}$. It turns out that this involution can be interpreted as a reflection of the statespace associated with the local algebra (§3.4.1). By stitching together these local statespace reflections with the help of the Haag-Kastler axioms and the Bisognano-Wichamnn theorem, we show that it is always possible to construct a global spatiotemporal reflection symmetry (§3.4.2). Since this symmetry stems from a total reflection of statespace, it doesn't just flip spatiotemporal degrees of freedom, it flips every discrete orientation structure associated with the theory, including charge (§3.4.3).

3.4.1 The Lie-Jordan Structure and Statespace Reflection

A C^*-algebra is an intrinsically oriented object. The involution, in conjunction with the associative C^*-product, naturally defines an orientation structure on the algebra. Intuitively, one can think of this structure as a way of keeping track of the order of multiplication between operators. After all, since multiplication is not commutative the order matters! Moreover, the presence of this privileged orientation allows the C^*-product to be decomposed into a pair of products, a Lie-Jordan structure, both of which play important roles in the physics.

As we saw in Ch. 2.1.3, a C^*-algebra is analogous to the complex plane, with the canonical involution implementing complex conjugation. The orientation structure on $\mathfrak{A}$ can be viewed as a choice of complex unit i. The defining feature of i is that it is a root of the equation $x^2 = 1$. Of course there are actually two roots, i and $-i$. While numerically distinct (since they are non-zero additive inverses), these numbers have identical algebraic properties. They are essentially mirror images of each other. We could select either as a complex unit and all of mathematics would turn out the same. Similarly, the positive and negative "imaginary halves" of the C^*-algebra, while distinct, have identical algebraic properties. In fact, algebraically it doesn't really matter that we picked left multiplication $(A, B) \mapsto AB$ over right multiplication $(A, B) \mapsto BA$ in the first place. Viewed from within the algebra, the order of multiplication matters, but from an external perspective, only the relative order is important.

This intuition can be formalized as follows: in the category of C^*-algebras, each algebra, $\mathfrak{A}$, has a unique dual, $\mathfrak{A}^{op}$, consisting of the same underlying linear space, involution, and norm, but equipped with opposite multiplication and opposite orientation. (If we let '$\cdot$' denote multiplication in $\mathfrak{A}$, we have $A \cdot B = B \cdot^{op} A$) The involution, itself a mapping from $\mathfrak{A} \to \mathfrak{A}$, can be used to define a natural isomorphism $\mathfrak{A} \to \mathfrak{A}^{op}$:

Fact 3.4. *The involution operation on $\mathfrak{A}$ naturally defines a $*$-isomorphism, $\varphi : \mathfrak{A} \to \mathfrak{A}^{op}$.*

In a strict mathematical sense, the involution is not a symmetry of $\mathfrak{A}$ since it is not an automorphism. It preserves all of the algebraic structure except for the order of products. Such a mapping is called an *anti-automorphism*. But here's the kicker: in physics we only use the self-adjoint elements of $\mathfrak{R}(\mathcal{O})$ to directly represent physical quantities. The non-self adjoint operators are used to represent algebraic relations (either between quantities or between certain quantities and certain symmetries). But such relations can be equally well encoded by the opposite choice of orientation, suggesting that this anti-automorphism should count as a symmetry of the local algebra.

To get a better sense of the role the orientation structure plays in physics, we can recast it in a slightly different form. The involution determines a canonical decomposition of the C^*-product into symmetric and antisymmetric components which preserve $\mathfrak{R}_{SA}(\mathcal{O})$:

$$AB = A \bullet B - i(A \star B), \tag{3.3}$$

The symmetric product, $\bullet$, is commutative but non-associative (i.e. $A\bullet(B\bullet C) \neq (A\bullet B)\bullet C$). It gives $\mathfrak{R}(\mathcal{O})$ the structure of a *Jordan algebra*. The antisymmetric product, $\star$, is associative but non-commutative. It gives $\mathfrak{R}(\mathcal{O})$ the structure of a *Lie algebra*. The presence of the involution thus reveals the original C^*-product to be two different multiplication operations disguised as one.

These new products are important because they play different representational roles in the physics. The Jordan product determines the spectrum of the self-adjoint operators and thus captures the role these operators play in representing measurable physical quantities. The Lie product enables us to define a natural exponential map which identifies elements of $\mathfrak{R}_{SA}(\mathcal{O})$ as the infinitesimal generators of connected spacetime symmetries. This captures the link between conserved quantities and spacetime symmetries given by Noether's first theorem.[31] Both are essential in quantum theory, hence both products are needed.[32]

Not only does the Lie product connect observables to symmetries, it also defines a natural geometric orientation on the (local) statespace associated with $\mathfrak{R}(\mathcal{O})$ — the collection of normalized, positive linear functionals over $\mathfrak{R}(\mathcal{O})$.[33] The involution acts anti-automorphically

[31]For example, the exponential map takes the energy-momentum operator to the generators of the translation subgroup of the Poincaré group, $P^\mu \mapsto e^{itP_\mu}$.

[32]The Jordan algebra program was a notable attempt to formulate quantum theory using only properties of the Jordan product. Although the program led to a number of advances in pure mathematics and quantum information theory, it was abandoned as a slimmed-down approach to quantum theory due to the lack of a coherent definition of the tensor product of two Jordan algebras, an essential component of any theory describing localized systems. In addition, general classification theorems showed that nearly all Jordan algebras which could be used in quantum theory are isomorphic to the symmetric part of a full C^*-algebra.

[33]The basic idea is as follows: the collection of states over $\mathfrak{R}$ has the structure of a convex set, the extremal points of which are the pure states. Each pair of pure states determines a face of the convex set, and each face is affinely isomorphic to either a straight line or a Euclidean 3-ball. Following Alfsen and Shultz, we define an *orientation* of the statespace to be a continuous choice of north and south poles for each facial 3-ball. If we represent the local statespace using a Hilbert space, the orientation corresponds to a choice

on this structure. While the natural isomorphism $\varphi : \mathfrak{A} \to \mathfrak{A}^{op}$ preserves the Jordan product, it reverses the Lie product, $A \star B \mapsto -(A \star B) = B \star A = A \star^{op} B$. In doing so it reverses the relationship between generators and symmetries and flips the orientation of the local statespace. The seeds of a PCT operator are contained in this local reflection symmetry. In the next section we'll examine how this symmetry can be extended to a spatiotemporal reflection of the entire net of local algebras. The key mathematical idea will be supplied by the Bisognano-Wichmann theorem (via the modular covariance).

3.4.2 Parity and Time Reversal on Wedges

Because of the Roberts-Roepstorff version of Wigner's theorem, we know that a global symmetry of the net must be implemented by a unitary or antiunitary operator. Our goal is to piece together local statespace reflections to generate a global symmetry. The first step is to extend these reflections to the local GNS Hilbert space. By the Reeh-Schlieder theorem, the vacuum is cyclic and separating for every local algebra, hence we can use Ω to define the conjugation operator S, mapping $A\Omega \mapsto A^*\Omega$ for all $A \in \mathfrak{R}(\mathcal{O})$. Moreover, as we'll see in Ch. 4, the Reeh-Schlieder theorem also shows that the local GNS Hibert space *just is* the global Hilbert space (Prop. 4.8). Unfortunately, in general S is not an isometry, but a reflection with an additional twist. Modular theory provides a way out. Taking the polar decomposition $S = J\Delta^{1/2}$, we know that the associated modular conjugation $J \in \mathfrak{R}(\mathcal{O})$ will always be an antiunitary operator.

Thus using modular theory, we can find a canonical antiunitary conjugation operator associated with each local algebra in the vacuum sector. We'll now turn our attention to wedge regions. The Bisognano-Wichmann relations indicate that ultimately it will be the wedge modular conjugation $J_{\mathcal{W}}$ which is interpreted as a PCT operator. Under the present assumptions, modular covariance gives the wedge modular unitaries, $\Delta_{\mathcal{W}}^{it}$, direct geometric

of complex structure on the Hilbert space. Unlike the manifold orientations discussed in §3.1, there are typically more than two possible orientations on statespace. See Alfsen and Shultz (2001) and Alfsen and Shultz (2003) for a comprehensive treatment.

significance — they generate wedge-preserving Lorentz boosts. In this section we'll show how $J_{\mathcal{W}}$ receives a corresponding spatiotemporal interpretation in light of its relationship to $\Delta_{\mathcal{W}}^{it}$.

From the Tomita-Takesaki theorem, we know that $J_{\mathcal{W}}$ defines a natural anti-isomorphism between $\mathfrak{R}(\mathcal{W})$ and its commutant $\mathfrak{R}(\mathcal{W})'$. Prima facie, we don't know what $\mathfrak{R}(\mathcal{W})'$ looks like. From the microcausality axiom, we know that the algebra associated with the spacelike complement of $\mathcal{W}$ must be included in it, i.e. $\mathfrak{R}(\mathcal{W}') \subseteq \mathfrak{R}(\mathcal{W})'$. But in principle, it could be much larger. It doesn't even have to be associated with a simply connected, open region at all. Fortunately, this is where modular covariance comes to the rescue.

Proposition 3.5. *Modular covariance implies wedge duality, namely* $\mathfrak{R}(\mathcal{W}') = \mathfrak{R}(\mathcal{W})'$.[34]

Since wedge duality holds, and we know that $J_{\mathcal{W}}\mathfrak{R}(\mathcal{W})J_{\mathcal{W}} = \mathfrak{R}(\mathcal{W})'$, it follows immediately that $J_{\mathcal{W}}\mathfrak{R}(\mathcal{W})J_{\mathcal{W}} = \mathfrak{R}(\mathcal{W}')$. So conjugation by J maps a given wedge algebra onto the opposite wedge algebra. This is where the particular geometry of the wedge becomes important. The opposite wedge is related to the original wedge by a spatial reflection across their common edge, hence the modular conjugation implements a reflection in one spatial degree of freedom, specifically the direction perpendicular to the edge of the wedge.[35]

But that's not all it does. The modular conjugation is also an antiunitary mapping, and therefore must reverse the Lie product:

[34] *Proof Sketch.* Let $(J_{\mathcal{W}'}, \Delta_{\mathcal{W}'})$, and $(J_{\mathcal{W}'}, \Delta_{\mathcal{W}'})$ denote the modular objects associated with $\mathfrak{R}(\mathcal{W}')$ and $\mathfrak{R}(\mathcal{W})'$ in the vacuum sector. From the main Tomita-Takesaki theorem, we have that $J_{\mathcal{W}}\Delta_{\mathcal{W}}J_{\mathcal{W}} = \Delta_{\mathcal{W}}^{-1} = \Delta_{\mathcal{W}'}$. Thus $\Delta_{\mathcal{W}'}^{it}$ acts geometrically on local algebras. In particular, modular covariance entails that

$$\Delta_{\mathcal{W}'}^{it}\mathfrak{R}(\mathcal{O})\Delta_{\mathcal{W}'}^{-it} = \Delta_{\mathcal{W}}^{-it}\mathfrak{R}(\mathcal{O})\Delta_{\mathcal{W}}^{it} = \mathfrak{R}(\Lambda_{\mathcal{W}}^{-1}\mathcal{O}) = \mathfrak{R}(\Lambda_{\mathcal{W}'}\mathcal{O}).$$

It immediately follows that the action of the modular automorphism group associated with $\mathfrak{R}(\mathcal{W})'$ leaves the subalgebra $\mathfrak{R}(\mathcal{W}')$ invariant. A theorem due to Takesaki shows that this condition is equivalent to the existence of a unique conditional expectation $\mathcal{E} : \mathfrak{R}(\mathcal{W})' \twoheadrightarrow \mathfrak{R}(\mathcal{W}')$ (see Takesaki (2000), Thm. 4.2). Since the vacuum state vector is cyclic for both $\mathfrak{R}(\mathcal{W})'$ and $\mathfrak{R}(\mathcal{W}')$, this conditional expectation is implemented by Longo's canonical endomorphism $\gamma : \mathfrak{R}(\mathcal{W}') \hookrightarrow \mathfrak{R}(\mathcal{W})'$ (see Longo (1987) for details). Since $\mathcal{E} = \gamma$ is bijective and $\mathcal{E}$ is the identity on $\mathfrak{R}(\mathcal{W}')$ by definition, it is the identity everywhere. So $\mathfrak{R}(\mathcal{W}') = \mathfrak{R}(\mathcal{W})'$. □

[35] This is the basic idea, although technically we still need to show that the adjoint action of $J_{\mathcal{W}}$ is a uniform reflection. Consider any wedge region $\mathcal{W}_1$ contained in $\mathcal{W}$. $\mathcal{W}_1 = \mathcal{W} + a$ where a is a translation vector in the direction orthogonal to the edge of $\mathcal{W}$. By modular covariance, the wedge unitaries act as wedge preseving boosts on $\mathcal{W}_1$, hence by the uniqueness of the polar decomposition, they are the modular unitaries for $\mathcal{W}_1$ too. It follows that $J_{\mathcal{W}} = J_{\mathcal{W}_1}$ and so $J_{\mathcal{W}}\mathfrak{R}(\mathcal{W}_1)J_{\mathcal{W}} = \mathfrak{R}(\mathcal{W}_1')$. Since Ad$J$ reflects any sub-wedge of onto its opposite wedge, AdJ must implement a uniform spatial reflection.

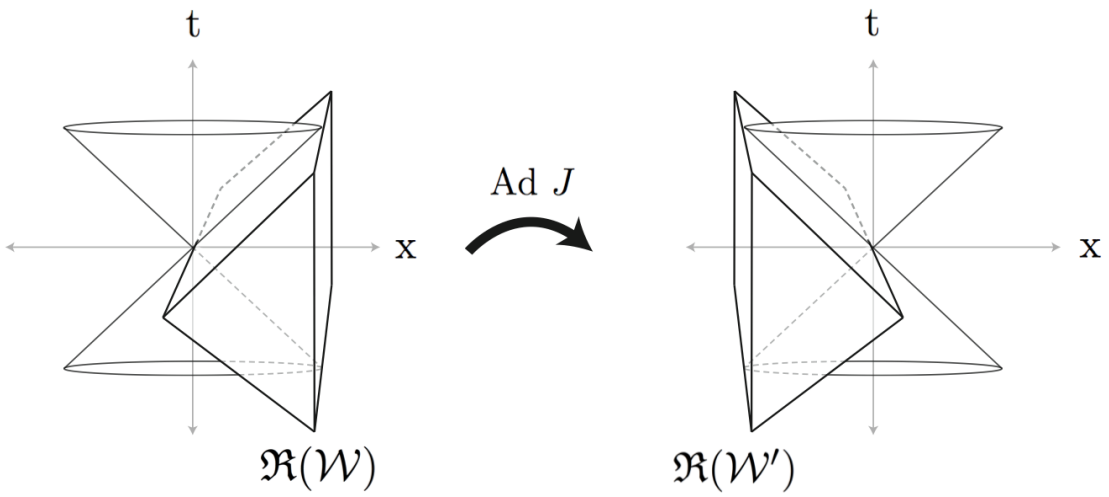


Figure 4: The P_1 reflection implemented by AdJ

Fact 3.6 (Alfsen-Shultz)**.** *A $*$-homomorphism between two von Neumann algebras preserves the Lie product iff it is unitarily implemented. It reverses the Lie product iff it is antiunitarily implemented.*[36]

As we have seen, the Lie structure encodes how operators act as generators of 1-parameter automorphism groups. In particular, the choice of lie structure fixes the infinitesimal generators of the translation subgroup:

$$\exp(t\delta_{iP^\mu})(x) = \exp(t(P^\mu \star x)) = e^{itP^\mu/2} x e^{-itP^\mu/2}. \tag{3.4}$$

Here P^μ, $\mu = 0, 1, 2, 3$, represent to the energy-momentum observables. The spectrum condition requires the spectral support of these generators to lie in the same momentum-space lightcone lobe. This in turn allows one to naturally define a temporal orientation on the wedge:

Proposition 3.7. *The Lie product along with the spectrum condition, naturally defines a temporal orientation, τ, on $\mathcal{W}$.*[37]

[36] For a proof, see Alfsen and Shultz (2001), Cor. 7.105.

[37] *Proof Sketch.* Recall from Appendix B that we employ the SNAG theorem to formulate the spectrum condition, requiring the support of the spectral measure μ to lie in a closed, proper cone $\overline{V} \in \hat{\mathcal{T}}$. This

Putting these two facts together, we see immediately that the modular conjugation not only maps $\mathcal{W}$ to $\mathcal{W}'$, but it also reverses the natural temporal orientation on the wedge. It does so by reflecting the generators of spatial motion (whose orientation are used to pick out a preferred direction of time), and thus also reverses motion as required of a time reversal symmetry operator.

Note that even though J is guaranteed to reverse time simply by dint of being an antiunitary transformation, it is only because it also has a well-behaved action on spatiotemporal regions that it counts as a kind of time reversal symmetry operator.[38] Strictly speaking a time reversal symmetry operator is a unitary/antiunitary operator that reverses the direction of time and motion and leaves everything else untouched. We can refer to more complicated symmetries like PT and PCT as kinds of time reversal symmetries because even though they change other quantities in the theory, they act on them in a uniform manner. Even though any antiunitary operator will reverse the natural temporal orientation, it may do so while completely scrambling the regions of spacetime. It's not even guaranteed to map open sets to open sets. Thus we see that the J operator is quite special. Its modular properties, along with wedge duality not only ensure that J implements a 1-direction spatial reflection, but also that it acts as a proper time reversal operator.

---

determines an orientation on energy-momentum space. We will show that this naturally determines a corresponding temporal orientation on Minkowski spacetime.

First note that the spectral measure μ is defined with respect to the complex structure on the vacuum Hilbert space carrying the unitary representation of $\mathcal{T}$. This complex structure is determined by the Lie product on the underlying algebra, thus the measure is unique up to a choice of the Lie product. (This connection can most easily be seen using the equivalent notion of a *Connes orientation*, a complex structure on the set of order derivations over $\mathfrak{R}_{SA}$. See Alfsen and Shultz (2003), Ch. 6 for details.)

Second, the theory treats $\mathcal{T}$ not as an abstract group, but as a group of translations on Minkowski spacetime, hence it comes equipped with a privileged embedding into Diff($\mathcal{M}$). The Minkowski metric determines a natural scalar product $(a, p) := a^0 p^0 - \vec{a} \cdot \vec{p}$ which defines an isomorphism between $\mathcal{T} \cong \hat{\mathcal{T}} \cong \mathbb{R}^4$. Thus $\overline{V} \in \hat{\mathcal{T}}$ is assigned a natural dual cone $V^d = \{a \in \mathcal{T} | (a, p) > 0, \text{ for } p \in \overline{V}\} \in \mathcal{T}$. Since μ is continuous and bounded, it is a tempered distribution on $\hat{\mathcal{T}}$, and so it has a well-defined Fourier transform. A theorem of Borchers (1996, Thm. II.1.7) shows that $\text{supp}(\mu) \in \overline{V}$ iff the Fourier transform $\mathcal{F}(\mu)$ is the boundary value of an analytic function holomorphic in the tube $T(\mathring{V}^d) = \{z \in \mathbb{C}^4 \mid \text{Im } z \in \mathring{V}^d\}$. Here, $\mathring{V}^d \in \mathcal{T}$ denotes the interior of the dual cone V^d. Via the embedding $\mathcal{T} \rightarrow$Diff($\mathcal{M}$), the spectrum condition therefore determines a natural temporal orientation on $\mathcal{M}$. □

[38] This point has recently been emphasized by Roberts (2013).

Using the covariance axiom and modular covariance, we can extend this action to a full "PT" operator acting on the net of observables.

Proposition 3.8. *The P_1T action of J is sufficient to construct a representation of the of the Poincaré group on the family of wedges. This includes a "PT" operator defined by* $\Theta := J_{\mathcal{W}} U(R_\pi)$.[39]

The scare quotes are included to emphasize that Θ may do more than just flip time and handedness. As we will see in the next section, it also conjugates charges making it a full PCT operator.

There is one important technical complication that can arise at this stage. Once we have an action of the full Poincaré group, we need show that this representation corresponds to the original representation given to us by the net covariance axiom. If we restrict attention to ordinary Bose/Fermi statistics and finite order parastatistics everything goes through fine. The case of infinite parastatistics poses a special problem since here one can have multiple inequivalent unitary representations of the Poincaré group present. Using this flexibilty, Streater has constructed counterexamples to the PCT and spin-statistics theorems for systems with infinite statistics.[40] In the Jost proof, these counterexamples are ruled out by the restriction to finite-component Wightman fields. In the present case, while there is no direct analogue, we can either directly restrict attention to cases of finite statistics, or impose a number of other conditions known to secure this result. The *split condition* is one natural possibility.[41] We'll discuss the split condition in more detail in Ch. 4.

[39] *Proof Sketch.* Ellers (2004) proves that the Lorentz group is bireflectional and that all involutions are conjugate. Using this fact, we can express any Lorentz transformation as the product of two P_1T wedge reflections. Following Buchholz and Summers (1993), we can generate a representation of the translations as well using P_1T reflections and modular covariance. Under the assumption of the split/nuclearity condition, these two representations give rise to a unique representation of the Poincaré group acting on the vacuum Hilbert space. □

[40] Streater (1967)

[41] Let $\mathfrak{R}(\mathcal{O}_1) \subset \mathfrak{R}(\mathcal{O}_2)$ be an inclusion of local algebras. Typically, the local algebras in any QFT will be type III_1 von Neumann factors. The split condition requires that there is a unique interpolating type I factor $\mathfrak{M}$ such that $\mathfrak{R}(\mathcal{O}_1) \subset \mathfrak{M} \subset \mathfrak{R}(\mathcal{O}_2)$. This allows us to import certain tools from standard quantum mechanics in which all algebras are type I. The split condition is implied by the nuclearity condition, which puts bounds on the number of local degrees of freedom in phasespace. Both conditions entail the compactness

3.4.3 Charge Conjugation

Probably the hardest part of the PCT puzzle is to understand how charge conjugation can be necessarily connected to a discrete spatiotemporal symmetry like PT. In a nutshell, the answer is that we have to look at how we are implementing the PT transformation in question. In effect we are performing the spatiotemporal reflection by flipping the Lie product structure, by changing how quantities and symmetries are linked at a fundamental level. As it turns out, the Lie structure is also important for defining the relationship between internal charges. In fact, we will show that flipping it systematically maps each charge to its conjugate.

Recall from §3.2 that localized charges can be represented by localized transportable endomorphisms. Conjugate charges have the defining property that $\varrho \circ \bar{\varrho}$ contains the identity. In the case of finite statistics, each charge has a unique conjugate (some charges may be self conjugate). Furthermore, one can show that

$$\bar{\varrho} = J_{\mathcal{W}} \varrho J_{\mathcal{W}}, \tag{3.5}$$

where $J_{\mathcal{W}}$ is the modular conjugation operator associated with any wedge.[42] Poincaré covariance immediately implies that $\bar{\varrho} = \Theta \varrho \Theta$.

The reason why this holds is somewhat subtle. From a mathematical angle, the category of localized transportable endomorphisms has a natural tensor product structure which is used for characterizing the composition of charges and hence the conjugacy equations. This tensor product depends on the collection of arrows in the category which is defined with respect to the C^*-product of the underlying algebra. The direction of the arrows is sensitive to the order of multiplication, hence is determined by the background Lie structure. An

of the global gauge group, an important assumption in DHR superselection theory. See Haag (1996), §V.5 for a discussion of how these three results are interrelated.

[42] For a proof, see Guido and Longo (1992).

arrow from $\varrho \to \varrho'$ is defined by

$$\mathrm{Hom}(\varrho, \varrho') = \{T \in \mathfrak{A} : T\varrho(A) = \varrho'(A)T, \forall A \in \mathfrak{A}\}. \tag{3.6}$$

Note that flipping the order of multiplication defines an arrow from $\varrho' \to \varrho$ in the original category. Since $J_{\mathcal{W}}$ reverses the Lie structure, it systematically reverses the direction of each arrow. Moreover, since $J_{\mathcal{W}}$ preserves localization, it takes objects (i.e. localized transportable endomorphisms) to objects. Conjugate objects can be characterized in terms of their arrows, they possess dual sets of arrows — for every arrow into ϱ from another object, ϑ, there's an opposite arrow from $\bar{\varrho}$ into ϑ. Thus any map like $J_{\mathcal{W}}$ which reverses arrows and acts automorphically on objects is a charge conjugation operator, sending $\varrho \mapsto \bar{\varrho}$.

Physically, we can think of the arrows in the category as encoding relations between charges. Conjugate charges are relationally defined; they have exactly opposite relations to other all charges. In this light, it is natural to view charge conjugation as a symmetry which reverses the network of arrows.

One of the greatest insights of DHR superselection theory is that this algebraic picture is equivalent to the more traditional picture of charge structure arising from conservation laws associated with internal gauge symmetries. If we have a net of observable algebras generated by Wightman fields, the observable net corresponds to the gauge invariant portion of the underlying field algebra. In this context, we can characterize superselection sectors using irreducible representations of the gauge group G. For instance in the Standard Model the gauge group is $SU(3) \times SU(2) \times U(1)$. The abelian $U(1)$ piece describes the symmetries of the electromagnetic force, while the non-abelian groups $SU(2)$ and $SU(3)$ characterize the symmetries of the weak and strong force respectively. The superselected charges correspond to conserved currents generated by these gauge transformations. The Wightman fields act on a a single, underlying Hilbert space H. Under the action of G, H splits into a direct sum of irreducible subspaces, $H = \bigoplus H_\sigma$, which are simultaneously irreducible representations

of the group G and unitarily inequivalent irreducible representations of the gauge invariant global algebra $\mathfrak{A}$. The original series of DHR papers prove that these representations are in 1-1 correpsondence with irreducible representations of $\mathfrak{A}$ fulfilling the DHR selection criteria. Conversely, the Doplicher-Roberts reconstruction theorem shows that given the category of DHR representations, one can naturally reconstruct a unique minimal field algebra and gauge group, G.[43]

In the Wightman field picture, we find that the action of $J_{\mathcal{W}}$ takes a given representation of G to its complex conjugate representation, representing the conjugate charge sector. We don't have to look far to see why. A representation (π, H) of G can be specified by a set of generators lying in the (weak closure) of the field algebra satisfying certain commutation relations, $[T^a, T^b] = if^{abc}T^c$ (where f^{abc} are the group structure constants). The complex conjugate representation is given by $-(T^a)^*$. Within these relations we immediately recognize the ubiquitous Lie product. The same structure which encodes how observables generate symmetries also encodes how unobservable field operators generate internal gauge transformations. Flipping this structure yields the complex conjugate representation. So at last the culprit behind the mysterious connection between C and PT symmetry is revealed. It is because the same Lie structure is employed in describing both internal and external symmetries that we find a systematic connection between them.

It is worth examining just how tight this connection is. Physicists have long suspected that in quantum theories, time reversal and charge conjugation are closely linked. For example, Bell (1955) argues that proper time reversal entails charge conjugation in much the same way that it entails velocity and momentum reversal. On his view, time reversal has to be a CT transformation. Others have argued that in one way or another the matter/antimatter distinction can only be made with respect to a background temporal orientation. Feynman (1949) famously proposes a picture of antiparticles as regular particles traveling backwards in time. Wallace (2009) contends that the distinction hinges on splitting the space of solutions

[43]See Halvorson and Müger (2006), §10 for a proof of the Doplicher-Roberts reconstruction theorem.

to the field equations into positive and negative frequency subspaces. This in turn depends on a choice of a complex structure associated with the direction of time.

Our discussion has revealed that there is a tight connection between charge structure and temporal structure in AQFT, but it turns out to be weaker than on any of the above views. The canonical Lie product in conjunction with the spectrum condition naturally defines a temporal orientation. This means that any proper time reversal transformation must be antiunitary. But the Lie structure also encodes the relational distinction between different charges, and since any time reversal symmetry must act uniformly on spatial degrees of freedom, it preserves charge localization. Thus it appears that any time reversal operator also conjugates charge.

There is an exception to this line of thought, though. As we saw, the relationship between conjugate charges is captured by the pattern of arrows in the category of localized transportable endomorphisms. One way to flip the direction of arrows systematically, and thus conjugate charges is by reversing the Lie product which appears in the definition of the relevant hom-sets. In theories where charges and anti-charges are treated symmetrically by the laws, however, it may be possible to define a unitary charge conjugation operation which reverses the direction of arrows while preserving the original Lie product. In this case one can combine an antiunitary CT reversal with the unitary charge conjugation procedure to produce a net antiunitary time reversal. This allows for theories like quantum electrodynamics which are invariant under P, C, and T symmetries independently.[44]

Thus, while T and C symmetries are closely related, this need not always be the case. They are not definitionally related like time and motion reversal. Indeed on the Feynman view it is hard to make sense of unitary charge conjugation at all. Since antiparticles just are particles traveling in the opposite temporal direction, any proper charge conjugation must involve a reverse of temporal orientation. The algebraic picture also has a distinct advantage over Wallace's view. The positive/negative frequency distinction he relies on only

[44]In some cases it may be necessary to use a unitary PC transformation, rendering the theory invariant under T, PC, and PCT transformations.

makes sense in the free field limit, where the theory can be linearized and solutions to the equations of motion can be expanded in terms of plane waves with opposite frequency. This leads Wallace to conclude that the matter/antimatter distinction is an emergent concept which only appears in the no-interaction limit where we can coherently talk about particle states. In contrast, the algebraic picture draws a fundamental distinction between matter and antimatter which is known to apply to interacting theories like 2-dimensional Yukawa theory.[45] After all, the PCT theorem is generally viewed as a deep, fundamental result. It would be odd if it turned out merely to describe high-level, emergent phenomena. The algebraic description of charge structure reinforces this idea.

3.4.4 Summary

We began with the definition of a quantum field theory as a net of von Neumann algebras over Minkowski spacetime modeling the Haag-Kaslter axioms. Every local algebra has a hidden reflection symmetry that reverses the algebra's natural Lie product. This product has two crucial roles. It specifies the generator link between observables and connected spacetime symmetries, including the translation subgroup, and hence in conjunction with the spectrum condition, naturally defines a temporal orientation on the net of algebras. It is also used to capture the relational difference between conjugate charges, appearing in the definition of arrows between localized transportable endomorphisms. Thus such a reflection reverses both temporal direction and conjugates charges.

Using tools from modular theory, including an additional assumption inspired by the Bisognano-Wichmann theorem called modular covariance, we showed that it was always possible to extend this local algebraic symmetry to a global symmetry of the net. The key was to recognize that the local reflection could always be decomposed into a positive piece Δ and an antiunitary reflection J. The action of the modular conjugation, J, on a local algebra maps the algebra onto its commutant. For wedge algebras, this means reflecting the wedge

[45] If a theory is asymptotically abelian and thus has a limiting particle interpretation, one can show how matter and antimatter are conventionally linked to opposite frequency wave solutions (Mund, 2001).

onto its spacelike complement, reversing a single spatial direction picked out by the wedge geometry. Not only does this show that the full symmetry operation must involve reflection of spatial degrees of freedom, it ensures that it preserves localization regions, a necessary condition to count as a genuine charge conjugation operator. Since the algebras associated with compact double-cone regions can be obtained by taking intersections of wedge algebras, this induces a P_1CT reflection symmetry on the entire net. Full PCT symmetry follows from the covariance axiom. (The split/nuclearity condition was required to ensure the uniqueness of the extended Poincaré action.)

Thus we find that three special features of J proved essential — (1) it is an antiunitary operator, (2) it is a modular conjugation operator, and (3) it preserves localization properties. Moreover, we see that it is really a mistake to view a PCT transformation as the product of three separate symmetry operations (as frequently assumed in Lagrangian proofs). Rather the transformation is a single operation that reverses the Lie bracket while preserving spatiotemporal localization properties. Sometimes it is possible to define unitary or antiunitary operators that reverse P, C, or T orientation individually, but this is not guaranteed. In contrast, modular theory is a generic feature of any AQFT. It is always possible to define a full PCT operator.

3.5 Greaves's Account

In her thought-provoking study of the PCT theorem, Greaves gives a very different account of the physical origins of PCT symmetry.[46] Her version of the tale has two main components. The first is a theory of antimatter based on Feynman's interpretation which analytically links time reversal and charge conjugation. If particles are traveling forwards and antiparticles backwards in time, if we reverse the direction of time, particles become antiparticles and vice versa. This collapses the PCT puzzle into a PT puzzle.

[46] Greaves (2010) is the primary source. This analysis is built upon work from her Ph.D. thesis, Greaves (2008). Arntzenius (2011) offers a similar account.

The second component of Greave's account is a PT theorem for classical field theories in Minkowski spacetime. The PCT theorem, she maintains, is essentially a relativistic phenomenon, or at least it has a direct analogue for classical field theories. The only apparent difference between her classical PT theorem and the quantum PCT theorem is the presence of antimatter which is easily handled by the Feynman picture. So according to Greaves any reasonable field theory, quantum or classical, has to be PT invariant, and if the theory has antiparticles then PT invariance is the same thing as PCT invariance.

There are two primary problems with Greaves's account. The first concerns the adequacy of the Feynman picture of antimatter. The second concerns the scope of the classical PT theorem. Taken at face value, the Feynman picture contradicts the axioms of algebraic QFT. After all, the spectrum condition requires that all particles have co-aligned 4-momentum. Everything has to be traveling in the same direction. Furthermore, as we have already seen, it is hard to make sense of unitary charge conjugation on this picture. As with all things Feynman, one should be skeptical of taking things at face value. A bit of digging is required to uncover the roots of the interpretation. The key lies in the physical gloss given to negative energy states that naturally appear as solutions to the Dirac equation in quantum electrodynamics. In the limit of no interactions, the theory admits an interpretation in terms of plane waves with opposite frequency. Conventionally, positive frequency solutions correspond to particle states and negative frequency solutions correspond to antiparticle states. This division into positive and negative frequency states requires a choice of complex structure on the theory's statespace (a special Hilbert space called a Fock space). The choice of complex structure is essentially a choice of a Lie product, and hence also picks out a preferred direction in time. As a result a link is forged between the direction of time and the frequency. Positive frequency solutions have a wavevector co-aligned with the direction of time and negative frequency solutions have an anti-aligned wavevector. Reversing the direction of time switches which wavevectors are co-aligned and anti-aligned. This just is the standard Lagrangian picture of antimatter defended by Wallace (2009). Feynman saw a

way of reading into this a story according to which annihilation reactions are just the same particle doubling back on itself in time. It simplifies certain problems, and has an elegant pictorial representation in terms of Feynman diagrams, but it is doubtful that he saw any deep metaphysical significance behind it.

Whether or not this is the view that Greaves has in mind is another question. She defends an interpretation inspired by the Feynman's view in her thesis and again in a paper with Frank Arntzenius.[47] In her version, particles are represented by oriented world-lines. Particles have world-lines co-aligned with the direction of time, anti-particles have anti-aligned world-lines. She remains open as to what exactly does the orientation work. It could be a 4-momentum. It could be a wavevector. It could be something else entirely. The problem is that quite a bit rides on what it is. If it's a wavevector as advocated above, then even if a time-reversal operation analytically entails that we relabel particles and anti-particles, it doesn't seem to say anything about *charge conjugation* as such. An electron will have -1 electric charge and a co-aligned wavevector. If we reverse the direction of time, its wavevector is now anti-aligned so we redescribe it as a positron. But intuitively, it should just be a positron with charge -1! If a positron is literally just an electron traveling backwards in time, then one could give the following kind of argument: from the perspective of the forward flow of time, the electron's path appears to be the motion reverse of what it is in the reverse direction. But since opposite charges affect motion in opposite ways, this is indistinguishable from a particle traveling forward in time with opposite charge. But this line of reasoning isn't open to the proponent of the wavevector view. Moreover, it seems to build a kind of (P)CT symmetry into the laws from the very start. Why should opposite charges affect (forward) motion in opposite ways?

Here is where the algebraic view does one better. The algebraic view makes it clear how opposite charges are characterized by opposite relational structures. This relational distinction is represented using the Lie product, the same structure which picks out a preferred

[47]See Greaves (2008), and Arntzenius and Greaves (2009).

direction of time via the spectrum condition. If we reverse time by changing the Lie product, by necessity we also wind up conjugating charges. Unitary charge conjugation is still permitted, since it is also possible to switch the direction of arrows without reversing the Lie product (in theories where charges and their conjugates are treated symmetrically). As discussed in §3.4.3, this picture also has the advantage over the wavevector view of making the distinction between matter and antimatter a fundamental rather than emergent one.

All this may just go to show that the best way to defend Greaves's analysis is to double down on the original intuition; anti-particles really are traveling in the opposite temporal direction from regular particles. While it is possible that one might be able to find an alternative set of axioms for QFT which replaces the spectrum condition with another stability assumption consistent with this picture of antimatter, it presents quite a tall order.

Even if this strategy proves successful, another difficulty looms. If particles and antiparticles are identical save for the orientation of their 4-momenta, then it is conceptually impossible for there to be particle/anti-particle pairs with different masses and spins. But this is theoretically possible. For example, Greenberg (2011) constructs a mathematically coherent QFT which violates the PCT theorem because the masses of particles differ from the masses of their partner anti-particles. The theory is not Lorentz invariant, however, so it fails to be a physically reasonable *relativistic* QFT.[48] It is important to note that the algebraic view does not build this restriction into its definition of antimatter. Instead antimatter is characterized solely in terms of its internal charge structure. This is all that is required for the possibility of creation/annihilation reactions. In Greenberg's theory, there are creation/annihilation events between particles and anti-particles with different masses, but because of the fact that charges are superselected, creation/annihilation interactions require particles and anti-particles to have exactly opposite charge. This is a clear advantage of the algebraic approach. On this view, it is the validity of the PCT theorem for Lorentz invariant QFTs which explains why particles and anti-particles in those theories have the

[48]It also violates the spin statistics connection, and Greenberg conjectures that it violates microcausality as well.

same mass, spin, and lifetimes. Indeed, from this vantage point, the PCT theorem explains why the Feynman interpretation is possible in the first place. It is because the theory is PCT invariant that we can interpret a forward moving antiparticle with charge $\overline{Q}$ as a backwards moving (opposite handed) regular particle with charge Q.

The second major problem concerns the scope of Greaves's account. Even if we accept her theory of antimatter, the classical PT theorem she relies on cannot recapture the full quantum PCT theorem. Her PT theorem considers classical field theories defined over Minkowski spacetime and has three central assumptions:

1. The fields transform as tensor quantities under diffeomorphisms of the spacetime manifold.

2. The laws have the form of partial differential equations (PDEs) which express the vanishing of some local polynomial combination of tensor fields.

3. The set of solutions is invariant under continuous Poincaré transformations.

The theorem shows that any such theory must also be invariant under a discrete PT transformation. The essential fact which Greaves's draws attention to is the following:

Fact 3.9. *In Minkowski spacetime it is impossible to directly encode* ***only*** *a temporal orientation using tensor fields. That is, any tensor field that is invariant under continuous Poincaré transformations is also invariant under T symmetry.*[49]

The construction of a temporal orientation we employed in §3.2 is ruled out because we made appeal to equivalence classes of tensor fields, and these, presumably are not the kinds of things that can appear in PDEs. Alternatively, we can pick out a temporal orientation with a timelike vector field, but this also defines a privileged spatial direction. Reversing the vector field amounts requires a P_1T transformation, which along with Lorentz covariance, entails invariance under full PT symmetry. If it were possible to directly encode just a temporal

[49]See Greaves (2010), fn. 12 for a proof.

orientation, then the theory would not be invariant under time reversal but invariant under parity inversion, hence fail to be PT invariant.[50]

As Greaves's notes, the situation is markedly different for field theories in non-relativistic Galilean spacetime, where it is possible to directly encode a temporal orientation using a special 1-form field.[51] In this setting, there is no analogue of the PT theorem to be found. The difference boils down to what kinds of orientation structures can be encoded by tensor fields on spacetime. In this sense, the PT theorem (and by extension the PCT theorem) is effectively a classical relativistic result according to Greaves.

The problem with this picture is that conditions (1) and (2) are not physically motivated. The requirement that the fields transform under a tensor representation of the Poincaré group excludes both pseudotensors and spinors. The former transform just like tensors under diffeomorphisms, but pick up a phase factor under discrete transformations. The latter are required for describing charged matter with half-integer spin, an essential ingredient in any theory like the Standard Model with a wide assortment of fermions in its particle zoo. Both kinds of fields can be used to construct counterexamples to the classical PT theorem, and yet both fall under the scope of the quantum PCT theorem.[52]

Similarly, it is difficult to see what motivates (2) apart from mathematical convenience. Generally covariant PDEs can take forms other than (2).[53] In addition, in AQFT, much of

[50]Greaves's theorem also requires that it is impossible to directly encode just a spatial orientation, a fact which she strangely does not discuss.

[51]We can represent a classical spacetime with $(\mathcal{M}, t_{ab}, h^{ab}, \nabla)$, where $\mathcal{M}$ is a 4-dimensional manifold, t_{ab} is a temporal metric, h^{ab} a spatial metric, and ∇ is a flat derivative operator on $\mathcal{M}$. In addition, we require that $h^{ab}t_{bc} = 0$, $\nabla_a t_{bc} = 0$, and $\nabla_a h^{bc} = 0$. A temporal orientation can be defined as a nowhere vanishing 1-form field, t_a, that satisfies the metric decomposition condition $t_{ab} = t_a t_b$ at every point in $\mathcal{M}$. See Malament (2012), §4.1 for more details.

[52]A static PT-pseudoscalar field is a particularly simple counterexample. Spinor fields can be used to construct PT-pseudotensors as well. The common bilinear currents $\bar{\psi}\gamma^\mu\gamma^5\psi$ and $\bar{\psi}\frac{i}{2}[\gamma^\mu, \gamma^\nu]\psi$ both pick up a sign under PT transformations, as does the partial derivative operator, ∂_μ. Here ψ and $\bar{\psi}$ are conjugate Dirac spinors and γ^μ are Dirac spin matrices. These pieces can be employed to construct PT violating Lagrangians. I explore the geometric relationship between spinor fields and orientation structure in Ch. 4 of my dissertation.

[53]Non-polynomial field theories are a rich area of study in mathematical physics. Here are three notable examples:

(i) The Sine-Gordon Lagrangian, $\mathcal{L}_{SG} = \frac{1}{2}(\varphi_t^2 - \varphi_x^2) - 1 + cos(\varphi)$

(ii) Liouville field theory, $\mathcal{L}_L = \frac{1}{4\pi} g_{\mu\nu}\partial_\mu\varphi\partial_\nu\varphi + e^{2b\varphi} + \frac{Q}{4\pi}R\varphi$

the interesting theoretical work is done by axioms which are not possible to formulate as PDEs at all. The spectrum condition is a useful example. We saw how in conjunction with the Lie product, it defines a natural temporal orientation. Interestingly, it does so without also defining an accompanying spatial orientation.

With regard to these problems of scope, the algebraic account sketched in §3.4 has a decisive advantage over Greaves's explanation. It expressly allows for the possibility of field operators transforming under non-tensor representations of $\mathcal{P}$. Moreover, it is extremely liberal with the mathematical form the laws can take. Orientation structures can be naturally defined by any mathematical means whatsoever. They do not have to be encoded directly by special tensor fields. To be fair, Greaves is well aware of the limited scope of her theorem. Still, she maintains that there is something puzzling about the PT connection even in the restricted class of classical tensor field theories. Presumably she thinks that the pattern of explanation will carry over directly to bosonic QFTs, which can be constructed without spinor fields. However, if we loosen her restrictions on laws and definability, it is doubtful that there will still be an analogue of Fact (5) available. More importantly, what makes the PCT theorem so interesting is that it holds for a broad, physically natural class of theories. The algebraic account supplies a common explanation which holds across this class. It is hard to see how Greaves's account could produce a similarly general explanation.

3.6 Classical or Quantum?

Unlike Greaves's analysis, the algebraic proof of the PCT theorem employs foundational assumptions from both relativity and quantum mechanics. Covariance and microcausality are essentially relativistic constraints, while the spectrum condition and modular theory are

(iii) Weinberg's chiral Lagrangian for π-mesons, $\mathcal{L}_\pi = (\partial_\mu \varphi)^2/(1 + f\varphi^2)^2$

The Sine-Gordon equation has soliton solutions and is used in research across particle physics, condensed matter theory, and optics. Liouville field theory has applications in 2-dimensional string theory, quantum gravity, and 3-dimensional classical gravity on negatively curved spacetimes. Weinberg used a quantized version of (iii) in his $SU(2) \times SU(2)$ model of π-mesons.

quantum mechanical in origin. This suggests that the full explanation of PCT symmetry requires ingredients from both theories. To what extent is the theorem unique to relativistic QFT?

In a classical field theory over Minkowski spacetime all observables have determinate values at all times, therefore the local observable algebras can be represented by commutative von Neumann algebras. But in a commutative algebra, there is no natural algebraic orientation since right and left mulitplication are indistinguishable. In addition, the modular structure becomes trivial. The modular automorphism group is trivial iff the generating state is tracial, but if the algebra is commutative, every state is a tracial state, so the modular flow reduces to the identity. As a result there is no meaningful analogue of the modular covariance condition. Furthermore, every abelian von Neumann algebra is *maximal*, i.e. $\mathfrak{M} = \mathfrak{M}'$.[54] Thus $J\mathfrak{M}J = \mathfrak{M}$, and so the modular conjugation operator cannot carry the geometric significance of a PCT operator. To compound these difficulties, Borchers has shown that no net of commutative von Neumann algebras can satisfy the spectrum condition.[55] And as we saw in the last section, counterexamples to Greaves's classical PT theorem can be generated either by loosening the restrictions on laws or by employing classical spinor or pseudo-tensor fields. Together, these results cast serious doubt on the existence of a classical theorem with similar generality and physical motivation as the quantum PCT theorem.[56]

Galilean QFT offers another useful test case. Greaves attempts to explain the lack of a PCT theorem in this domain by pointing out that in Galilean spacetime (unlike Minkowski spacetime), it is possible to use tensor fields to directly encode a temporal orientation without

[54]To see why, note that since $\mathfrak{M}$ is abelian, $\mathfrak{M} \subseteq \mathfrak{M}'$. Moreover, any vector cyclic for $\mathfrak{M}$ will also be cyclic for $\mathfrak{M}'$, therefore $\mathfrak{M}' = J\mathfrak{M}J \subseteq J\mathfrak{M}'J = J(J\mathfrak{M}J)J = \mathfrak{M}$.

[55]Borchers (1996), Thm. IV.6.2. Actually Borchers proves something even stronger, no net of commutative C^*-algebras can satisfy the spectrum condition.

[56]As we'll discuss in Ch. 5.3, there are reasons to believe that the appropriate classical analogue of a noncommutative C^*-algebra is not a commutative C^*-algebra but a Poisson algebra which has an associative Jordan product and a Lie product determined by the Poisson bracket. Weinstein (1997) has proposed an analogue of the modular automorphism group in this setting. Whether or not this can yield a classical analogue of the PCT theorem remains to be seen. At this stage no analogue for the modular conjugation is known. Additionally, in a noncommutative C^*-algebra, the Jordan and Lie products are more closely related to one another since they are both determined by the decomposition of the C^*-product.

any excess baggage. But as we've seen, this is an artificial restriction on definability. In principle though, the explanation must boil down to a difference between the structure of Galilean and Minkowski spacetime. One might hold out hope that orientation structures in Galilean spacetime are even more independent than in Minkowski spacetime. After all, the causal structure of Minkowski spacetime is encoded in a single spatiotemporal metric, while in Galilean spacetime there are two orthogonal metrics representing spatial and temporal distances separately. Although it is an attractive possibility, it turns out to be false. Once we define the Galilean analogues of temporal, spatial, and total orientation, we find that any two orientation structures naturally define the third, just as in Minkowski spacetime.[57]

The difference seems to come from the fact that the causal structure of Galilean spacetime is weaker than the causal structure of Minkowski spacetime. In relativistic theories, we require matter fields to propagate inside light cones. Signals cannot travel faster than c. In contrast, Galilean spacetime has no built-in speed limit. This is reflected in the fact that the Galilean version of microcausality is a significantly weaker condition. Observables on a given simultaneity slice are required to commute with each other, but this places no constraints on commutation relations at subsequent times. Relativistic microcausality is a much more potent, 4-dimensional condition. Galilean microcausality only applies to a 3-dimensional submanifold of spacetime.

This is not the only significant difference. Unlike in the Poincaré group, in the Galilei group temporal and spatial translations are distinct subgroups and thus have distinct generators in the corresponding Lie algebra.[58] This reflects a deep sense in which time and space are separate entities in Galilean spacetime and leads to substantial modifications of the spectrum condition. In Galilean QFT, the stability of matter is secured by a mass superselection rule along with a lower bound imposed on the spectrum of a particle's internal

[57] See Weatherall (2011), Appendix B for a discussion of this point.

[58] Recall that the Poincaré group is the semidirect product of the spatiotemporal translations with the Lorentz group, $\mathcal{P} = \mathcal{T} \rtimes \mathcal{L}$. The Galilei group is significantly more complex: $\mathcal{G} = \mathcal{U} \rtimes (\mathcal{S} \rtimes \mathcal{R})$, where $\mathcal{U}$ includes both spatial translations and Galilei boosts, $\mathcal{S}$ are the time translations, and $\mathcal{R}$ are the rotations.

energy (which as it turns out is a completely free parameter in the theory).[59] These restrictions are insufficient to endow either the modular operators or Wightman functions with the analyticity properties needed to prove the PCT theorem.

This discussion indicates that the spectrum condition is not entirely a quantum mechanical axiom. While it ensures stability, a quantum mechanical constraint, it does so by imposing restrictions on the energy-momentum operators, whose character is sensitive to the background spacetime structure. Relativistic and quantum considerations are more entangled in the axioms for AQFT than it initially appears. What is clear, however, is that the algebraic proof of the PCT theorem hinges on subtle interplay between the spectrum condition and microcausality. In classical field theories we can't even formulate the spectrum condition while in Galilean QFT, the analogues of microcausality and the spectrum condition are to weak to secure the result. While more work is needed to fully clarify the role of these axioms in AQFT, our preliminary conclusion is that, contra Greaves, the PCT theorem essentially depends on the unification of relativity and quantum mechanics.[60]

This conclusion is broadly consistent with the historical development of QFT. Initial attempts to relativize quantum theory in the late 1920s ran aground on a cluster of problems stemming from a deep conflict between relativistic causality and energy positivity. At the time, physicists were searching for a Lorentz invariant wave equation which could play the

[59] Lévy-Leblond (1967)

[60] Greaves and Thomas (2012) claim to give a proof of a classical PCT theorem drawing upon a more general result dubbed *strong reflection invariance*. There are direct parallels between strong reflection invariance and a global reversal of the Lie product, thus the relationship between these new results and the account I've given here should be explored further. Regarding the issue of classical PCT symmetry, however, it must be emphasized that their classical PCT theorem is markedly different from its quantum counterpart. True, they are both instances of the same theorem schema (Thm. 9.6):

> *If a classical/quantum formal field theory is (a) supercommutative and (b) invariant under a representation of the (cover of the) connected Lorentz group, then the theory is invariant under PCT reflections iff it is \$-hermitian.*

The problem is that different notions of \$-hermiticity are required to prove the classical and quantum versions. In the quantum case, the relevant notion is defined with respect to the canonical involution. In the classical case, it is defined with respect to the charge conjugation involution. Thus the classical theorem asserts that if a classical field theory obeys the spin-statistics connection and is invariant under a (cover) of the Lorentz group, then it is CPT invariant iff it is C invariant. The quantum PCT theorem permits C violating theories which are nonetheless still PCT invariant! Thus there is interesting interaction between spatiotemporal and charge structure in the quantum theorem which is entirely absent in the classical version.

role of the Schrödinger equation in standard quantum mechanics. The most direct strategy begins with the Schrödinger equation and replaces the the classical energy dispersion relation $E = \mathbf{p}^2/2m$ with the relativistic dispersion relation $E = \sqrt{\mathbf{p}^2 + m^2}$. While the resulting equation is Lorentz invariant, it has a number of problems including the following — if we begin with a wavepacket, ψ, describing a particle localized in some region, as the particle evolves through time, the wavepacket spreads out faster than c. As a result the probability of detecting the particle outside its own lightcone is nonzero. This raises the specter of faster-than-light signaling and other causality-violating paradoxes. Special relativity requires something more than just Lorentz invariance. It requires *hyperbolic* equations of motion describing localized initial data whose time evolution propagates with speed less than c.

One can avoid this problem by starting with a hyperbolic wave equation like the Klein-Gordon or Dirac equation. Unfortunately, such equations typically have both positive and negative energy solutions. The presence of unbounded negative energy solutions mean that it is always possible for a system to decay into lower and lower energy states. There will be no stable matter in such a theory. Furthermore, the negative energy solutions create serious difficulties for interpreting $|\psi|^2$ as a probability density, destroying any hope for recovering the probabilistic formalism of standard quantum mechanics.[61] One can implement a cutoff condition which throws away the negative energy states, but only at the cost of ruining hyperbolicity or Lorentz invariance.

These obstacles can be turned into a rigorous no-go theorem:

Theorem 3.10 (Strocchi)**.** *There is no wave equation describing a finite, fixed number of quantum particles which is simultaneously Lorentz invariant, hyperbolic, and energy-positive.*[62]

[61] In non-relativistic quantum theory, $|\psi|^2$ is interpreted as a probability density, ρ, satisfying a continuity equation of the form $\partial\rho/\partial t + \nabla \cdot \mathbf{J} = 0$. But this equation is not relativistically invariant, leading to non-conservation of the associated probability current $\mathbf{J}$. Order can be restored by replacing $\mathbf{J}$ with a relativistically invariant current, j_μ, but now the time component, j_0, is no longer positive definite, yielding negative probabilities. (Interestingly, the Dirac equation does not suffer from this problem, and was therefore an important stepping stone towards QFT.) See Strocchi (2013), Ch. 1.2 for a discussion of this problem.

[62] See Strocchi (2013), Prop. 2.2.

QFT effectively sidesteps this problem by dropping the requirement that the theory describe a finite, fixed number of particles. Indeed, in AQFT the theorem's three primary restrictions reappear as axioms — covariance, microcausality, and the spectrum condition, respectively. The negative energy states never really go away. Rather they are reinterpreted as positive energy states with opposite charge. Antiparticles and creation/annihilation processes thus emerge as critical ingredients in building a consistent QFT. But not just any old antiparticles will do the trick. They need to have exactly the same mass and spin as their conjugate partners. A number of partial results indicate that upsetting this balance will lead to violations of microcausality, Lorentz invariance, or both.[63] Viewed from this angle, PCT symmetry plays a fundamental explanatory role in QFT. It is only because the theory is PCT invariant that we can reinterpret negative energy states as describing antiparticles in a manner consistent with the requirements of relativistic causality.

[63]See Greenberg (2002), (Peskin and Schroeder, 1995, §2.4), and Weinberg (2005, §5.1) for different formulations of this kind of argument. To date, none of them have been made fully rigorous.

Chapter 4

The Type Question

"QFT is the theory of local observables and non-local states." — folklore

What does it mean to say that a state is localized in a region, $\mathcal{O}$? Intuitively, a localized state represents something interesting going on inside $\mathcal{O}$ and nothing everywhere else. Of course in QFT "nothing" is really quite interesting. By the Reeh-Schlieder theorem, the vacuum state is cyclic and separating for every local algebra. The latter condition entails that any non-zero projection operator has a non-vanishing expectation value, $\omega(E) > 0$. Thus any quantum event has some chance of spontaneously occurring in the vacuum! For an arbitrary local observable, fluctuations around the expectation value $\omega(A)$ are quantified by the variance $\omega(A^2) - \omega(A)^2$. Since ω is represented by a vector state, Ω, in the global GNS Hilbert space, the variance vanishes iff Ω is an eigenvector for A. But if this is true, then $(A - \lambda I)\Omega = 0$, and since Ω is separating, $A = \lambda I$. So any non-trivial observable has positive variance. The QFT vacuum is not "nothing." It is a roiling sea of fluctuations.

Some care is therefore needed in order to precisify the idea of a localized state. The most widely adopted definition is due to Knight (1961). A state, ϕ, is *Knight-localized* in $\mathcal{O}$ whenever

$$\phi(A) = \omega(A), \quad \forall A \in \mathfrak{R}(O').$$

The requirement that ϕ and ω have the same expectation values for observables in the spacelike complement $\mathcal{O}'$ captures the intuition that a localized state "looks like" the vacuum outside of $\mathcal{O}$.[1]

It is tempting to assume that we can generate localized states simply by applying local operators from $\mathfrak{R}(\mathcal{O})$ to the vacuum. The cyclicity of Ω thwarts this idea. Since $\overline{\mathfrak{R}(\mathcal{O})\Omega} = \mathcal{H}$, we can approximate any global state to arbitrary accuracy by applying polynomials of local observables to Ω. Consequently, a typical state represented by $\Psi = A\Omega$, $A \in \mathfrak{R}(\mathcal{O})$, will not be Knight-localized in $\mathcal{O}$. If $||A|| = 1$, however, it follows from the Reeh-Schlieder theorem that A must be an isometry (i.e. $A^*A = I$), and Ψ will be Knight-localized in $\mathcal{O}$. As $||A||$ increases, Ψ becomes increasingly delocalized, thus we can use the algebraic norm as a rough measure of the degree of delocalization of Ψ. Interestingly, the set of Knight-localized states span the entire global Hilbert space. Hence we can also approximate any global state by taking linear combinations of localized states.[2]

Physically, this is telling us that the vacuum state (or more generally any state analytic for the energy) is entangled across all length scales. There are non-trivial correlations between observables in $\mathfrak{R}(\mathcal{O})$ and every other region. Even though these correlations decay rapidly with distance, they never vanish. Thus by judiciously choosing operations from $\mathfrak{R}(\mathcal{O})$, we can exploit the information contained in these correlations to replicate any other state, including states localized in regions spacelike separated from $\mathcal{O}$.[3]

[1] It is important not to conflate the concept of a *local state* with that of a *localized state*. A local state is the restriction of a global state to a local algebra of observables. A Localized state is a global state which looks like the vacuum outside of a particular region.

[2] This follows from the fact that any element of a C^*-algebra is a finite linear combination of unitary elements (Kadison and Ringrose, 1997a, Thm. 4.1.7) and the observation that for any unitary, $U\Omega$ represents a Knight-localized state.

[3] Taken at face value, this seems to tell us that by performing local measurements in $\mathcal{O}$ we can manipulate the state in far away regions. If truly possible, this would be a gross violation of relativistic causality. Clifton and Halvorson (2001) argue that the key to resolving this paradox lies in the fact that the measurement procedures required to change the global state in this fashion will be (by necessity) selective. Any local observer hoping to control events across the universe only has non-selective measurement procedures at her disposal. Even though unitary operations are non-selective, arbitrary linear combinations of unitary operations are not. This explains why it is possible to use localized states of the form $U\Omega$ to approximate any global state.

These ubiquitous non-local correlations leave their mark on the intrinsic algebraic structure of the local observables. A series of results culminating with Fredenhagen (1985a), establish that in virtually all physically salient theories the local von Neumann algebras are isomorphic to the unique hyperfinite type III_1 factor. This stands in stark contrast to non-relativistic QM (where algebras are typically finite type I_n factors), and has important ramifications for entanglement, localization, and the treatment of measurement in QFT. Recall that for a type III_1 algebra, the state space orbit diameter $d(\mathfrak{R}) = 0$. Thus the kinematic unification of relativity and QM appears to require that the observables assigned to spatiotemporal regions are maximally noncommutative. As a consequence, entanglement is the norm rather than the exception and it is extremely difficult to isolate spatiotemporal subsystems from the environment.

Modular theory is intimately related to the type question. As we have already seen, the spectral features of the modular operators give rise to the Connes classification of type III von Neumann algebras. In this chapter (§4.1), we discuss how the scaling behavior of the local modular automorphism groups establish the type III property. Thus we find that the modular structure of the local algebras lies behind many of the stranger features of relativistic quantum theory.

Despite this, all is not lost. While at a fundamental level the quantum field is a highly entangled mess, at an emergent level, there can be well-behaved localized structure. The standard Haag-Kastler axioms do not guarantee this, however, prompting the search for additional axioms. Once again we find modular theory playing a central role. In §4.2 we'll examine two candidate axioms, *nuclearity* and the *split condition*. Although initially inspired by physical arguments, both have surprising mathematical connections to modular theory. We'll go on to explore how split localization can be used to define local charge observables and a local Hamiltonian operator.

The final section (§4.3), contains a critique of a new ontological picture of quantum theory advanced by Wallace and Timpson called *spacetime state realism*. According to this view,

the fundamental ontology of QFT consists of a field of local density operators indexed by spacetime region. As we'll see, it is difficult if not impossible to make sense of this picture when the local algebras are type III.[4]

4.1 The Type III Property

The aim of this section is to clarify how the type III property is related to entanglement and to discuss the physical ideas behind the proof of the type III property. Following Clifton and Halvorson (2001), we will adopt a strong definition of entanglement. A global state, ϕ, is entangled across regions $\mathcal{O}_1$ and $\mathcal{O}_2$ if the restriction of ϕ to the joint algebra $\mathfrak{R}(\mathcal{O}_1) \vee \mathfrak{R}(\mathcal{O}_2)$ is not contained in the ultraweak closure of the convex hull of product states on the joint algebra. A *product state* on $\mathfrak{R}(\mathcal{O}_1) \vee \mathfrak{R}(\mathcal{O}_2)$ is a state, ρ, for which there exist local states, ρ_1, ρ_2, on $\mathfrak{R}(\mathcal{O}_1)$,$\mathfrak{R}(\mathcal{O}_2)$ such that $\rho(AB) = \rho_1(A)\rho_2(B)$ for all $A \in \mathfrak{R}(\mathcal{O}_1)$, $B \in \mathfrak{R}(\mathcal{O}_2)$. Product states only exhibit classical correlations. According to the definition above, an entangled state cannot be approximated by ultraweak sequences of convex combinations of product states. Therefore entangled states exhibit correlations which depart significantly from classical statistics.[5]

The Reeh-Schlieder theorem entails that the any state analytic for the energy is entangled across all local algebras. As a direct result, any physically reasonable QFT possesses a dense set of entangled states (Summers, 1990). In this sense entanglement is generic in QFT. If duality holds, $\mathfrak{R}(\mathcal{D})' = \mathfrak{R}(\mathcal{D}')$, and we can represent the quantities characterizing the spacelike environment relative to a region using the commutant of the local algebra. The anti-isomorphism $J\mathfrak{R}(\mathcal{D})J = \mathfrak{R}(\mathcal{D}')$ induced by the modular conjugation can then be viewed as reflecting a corresponding structural equivalence between the local and environmental observables. If the local algebras are type III factors, then an even stronger result follows: there are no product states across $\mathfrak{R}(\mathcal{O}) \vee \mathfrak{R}(\mathcal{O})'$. This essentially follows from the fact that

[4]Material from this chapter was originally presented at a colloquium at UC Irvine in May 2014.

[5]Note that $\mathfrak{R}(\mathcal{O}_1) \vee \mathfrak{R}(\mathcal{O}_2)$ will have entangled states iff both $\mathfrak{R}(\mathcal{O}_1)$ and $\mathfrak{R}(\mathcal{O}_2)$ are nonabelian.

$\mathfrak{R}(\mathcal{O}) \vee \mathfrak{R}(\mathcal{O})'$ cannot be written as a tensor product . (Since $\mathfrak{R}(\mathcal{O})$ is a factor, $\mathfrak{R}(\mathcal{O}) \vee \mathfrak{R}(\mathcal{O})' = B(\mathcal{H})$ which is type I, but the tensor product $\mathfrak{R}(\mathcal{O}) \otimes \mathfrak{R}(\mathcal{O})'$ is type III (Kadison and Ringrose, 1997b, Prop. 11.2.26).) In this case the type III property tells us that the quantum field inside a double cone can never be disentangled from the field in its spacelike complement.[6]

If the local algebras were instead type I factors, then it would be possible to find product states across the field in $\mathcal{O}$ and $\mathcal{O}'$. Given an arbitrary type I factor, $\mathfrak{N}$, acting on a Hilbert space, $\mathcal{H}$, there is an isomorphism $\mathcal{H} \cong \mathcal{K}_1 \otimes \mathcal{K}_2$, where $\mathcal{K}_1$, $\mathcal{K}_2$ are Hilbert spaces such that $\mathfrak{N} \cong B(\mathcal{K}_1) \otimes I$ and $\mathfrak{N}' \cong I \otimes B(\mathcal{K}_2)$. Therefore we can write $B(\mathcal{H}) = \mathfrak{N} \vee \mathfrak{N}' \cong \mathfrak{N} \otimes \mathfrak{N}'$. In non-relativistic QM, this result provides a standard recipe for representing composite systems. The type III property signals a breakdown of this tidy picture.[7]

Locally, this highly entangled state of affairs manifests itself in the fact that there are no pure, normal states on the local algebras.[8] Every state of $\mathfrak{R}(\mathcal{O})$ which can be represented by a density operator is mixed. Moreover, every state is *intrinsically mixed*, i.e. disjoint from every pure state on $\mathfrak{R}(\mathcal{O})$. If a state is intrinsically mixed, then its GNS representation is reducible, every subrepresentation is reducible, every subrepresentation of every subrepresentation is reducible, etc. ad infinitum. It follows that every component of a local state is itself mixed, so we cannot interpret the states as a mixture of pure states.

Additionally, the type III property entails that for every projection $E \in \mathfrak{R}(\mathcal{O})$ there is an isometry $V \in \mathfrak{R}(\mathcal{O})$ such that

$$V^*V = I, \quad VV^* = E. \tag{4.1}$$

[6]See Clifton and Halvorson (2001).

[7]In §4.3 we'll see how this spells trouble for spacetime state realism.

[8]Note that this holds for any type III or type II von Neumann algebra in virtue of the fact that they lack minimal projections. A local state ρ is pure iff the kernel $\mathcal{L}_\rho = \{A \in \mathfrak{R}(\mathcal{O}) \mid \rho(A^*A) = 0\}$ is a maximal left ideal (Kadison and Ringrose, 1997b, Thm. 10.2.10.). Every normal state has a unique support projection, E_ρ, defined as the meet of all projections $E \in \mathfrak{R}(\mathcal{O})$, such that $\rho(E) = 1$. Since $\mathfrak{R}(\mathcal{O})$ lacks minimal projections by hypothesis, there is some nonzero projection $F < E_\rho$. Since $\mathfrak{R}(\mathcal{O})$ has a separating vector (by the Reeh-Schlieder theorem), every state is a vector state and so there is some state ψ represented by a vector which lies in the range of F. It follows that $\mathcal{L}_\rho \subset \mathcal{L}_\psi$, so $\mathcal{L}_\rho$ isn't maximal.

Thus every projection is equivalent to the identity, $E \sim I$. This distinguishes the type III case from the type II case, where every non-zero projection E has a subprojection $F < E$ that is finite (i.e. for any $F' < F$, if $V^*V = F$ and $VV^* = F'$, then $F' = F$).

As a result of this fact, Yngvason (2005) shows that we cannot use positive operators to measure local deviations from the vacuum. Given a global state we can quantify how much it deviates from the vacuum in a region by taking

$$D_\phi(\mathcal{O}) \equiv \sup\{|\phi(A) - \omega(A)| \mid A \in \mathfrak{R}(\mathcal{O}),\ ||A|| \leq 1\}. \tag{4.2}$$

If the local algebras are type I, there exist positive operators, $P_\phi \in B(\mathcal{H})$ such that $D_\phi(\mathcal{O}) = 0$ iff $\omega(P_\phi) = 0$, for a dense set of states $\{\phi\}$. (Such operators cannot be local observables, however, due to the Reeh-Schlieder theorem.) For type III algebras, no such operators exist for any state. Buchholz and Yngvason (1994) go on to argue that importing illicit assumptions from the type I case to the realm of QFT underwrites a number of apparent causality-violating paradoxes.[9]

Thus the type III property has far reaching implications for the foundational structure of QFT. Following partial results obtained by Kadison (1963) and Guenin and Misra (1963), Araki (1964b) proved that the local algebras in free field theories are type III factors. This result entails that any theory which is *locally quasiequivalent* to a free theory must also have the type III property. Two nets of von Neumann algebras $\mathfrak{R}_1$, and $\mathfrak{R}_2$ are locally quasiequivalent if for each double cone there exists an isomorphism $\varphi : \mathfrak{R}_1(\mathcal{D}) \rightarrow \mathfrak{R}_2(\mathcal{D})$. When one of these nets corresponds to a free field theory (i.e. the Fock representation) the other net corresponds to a theory which is said to be *locally normal* (Haag, 1996, Ch. III.3.1). Concrete examples of such theories include 2-d Yukawa theory and 2-d scalar theories with polynomial interaction terms.

[9]The Fermi two-atom problem describes a spacelike separated pair of atoms, one of which is in its ground state and one of which is in an excited state. Hegerfeldt (1994) argues that if the second atom spontaneously decays, then the emitted radiation has a non-zero probability of exciting the first atom instantaneously. Buchholz and Yngvason show that Hegerfeldt's argument relies on the existence of certain minimal projection operators which cannot be elements of a local type III algebra.

While every known example of a theory satisfying the Haag-Kastler axioms has the type III property, at this stage we still lack a general proof from the standard axioms (1)-(6). The closest result is an important lemma by Borchers (1967, Thm. III.3.b), which shows that for any region strictly contained in another, every projection from the smaller region is infinite relative to the larger's observables:

Theorem 4.1 (Borchers)**.** *Assume Haag-Kastler axioms (1)-(6). If* $\mathcal{O}_1 \ll \mathcal{O}_2$ *then every projection in* $\mathfrak{R}(\mathcal{O}_1)$ *is infinite relative to* $\mathfrak{R}(\mathcal{O}_2)$*.*

For unbounded regions with particular geometry, more definite results are known. A particularly important example are wedge algebras. Work by Driessler (1977) and Longo (1979) using modular constructions called *half-sided translations*[10] has established that the algebras for wedges are type III_1:

Theorem 4.2 (Driessler-Longo)**.** *Assume Haag-Kastler axioms (1)-(6). For any wedge* $\mathcal{W}$*, the local algebra* $\mathfrak{R}(\mathcal{W})$ *is type* III_1*.*[11]

The proof of Thm. 4.2 relies on the particular geometry of the wedge and therefore cannot be applied to compact regions. Fredenhagen (1985a) realized that one can exploit a different piece of modular data, the Bisognano-Wichmann relations, to establish the type III property for double cones.

[10]Given a von Neuman algebra with a cyclic, separating vector, a half-sided translation is a continuous, positively generated group of unitary operators, $U(s)$, $s \in \mathbb{R}$, that leaves the cyclic, separating vector invariant and acts endomorphically on the algebra (i.e. $U(s)\mathfrak{M}U(-s) \subset \mathfrak{M}$). Whenever such translations exist, they are intimately related to the modular objects associated with the region:

$$\Delta^{it}U(s)\Delta^{-it} = U(e^{-\pi t}s)$$
$$JU(s)J = U(-s).$$

Due to the spectrum condition and the geometry of the wedge, the subgroup of spacelike translations in the direction picked out by the wedge's characteristic two-plane are half-sided translations for $\mathfrak{R}(\mathcal{W})$.

[11]The wedge algebra is type III_1 iff the modular automorphism σ^t is outer for all $t \neq 0$. Assuming for reductio that σ^t is inner for at least one t, one can show that the unitary operators $\sigma^t(A) = WAW^*$ must belong to the centralizer of $\mathfrak{R}(\mathcal{W})$. But this is not consistent with the existence of half-sided modular translations, since it follows from their relationship with the modular objects that $U(s)E\Omega = E\Omega$ for any projection E in the centralizer. Moreover, $\mathbb{C}\Omega$ is the unique invariant subspace under the half-sided translations, thus $E\Omega = \Omega$. Since Ω is separating, it follows that either $E = 0$ or $E = I$. So $\mathfrak{R}(\mathcal{W})$ is either trivial or type III_1.

Fredenhagen's idea was to imbed a double-cone region in the corner of a wedge, tangent to the wedge at the origin. In the short distance scaling limit, one can show that the modular automorphisms of $\mathfrak{R}(\mathcal{D})$ and $\mathfrak{R}(\mathcal{W})$ approximately coincide. The Bisognano-Wichmann theorem entails that the modular objects associated with $\mathfrak{R}(\mathcal{W})$ in the vacuum sector implement geometric transformations of the net. In particular, $\Delta_{\mathcal{W}}^{it} = U(\Lambda_{\mathcal{W}}(2\pi t))$. From the covariance axiom, it follows that $\Delta_{\mathcal{W}}^{it}\mathfrak{R}(\lambda\mathcal{D})\Delta_{\mathcal{W}}^{-it} \subset \mathfrak{R}(\mathcal{D})$ for $|t| \leq (2\pi)^{-1}|\ln\lambda|$, $\lambda < 1$. This fact can be used to obtain the following estimate:

Lemma 4.3 (Fredenhagen)**.** *For each $f \in L^1(\mathbb{R})$ and $0 < \lambda < 1$, there exists a constant $c_f(\lambda) > 0$ such that $c_f(\lambda) \to 0$ as $\lambda \to 0$, and*

$$||\int dt\ f(t)\{\Delta_{\mathcal{D}}^{-it} - \Delta_{\mathcal{W}}^{-it}\}A\Omega||^2 \leq c_f(\lambda)\{||A\Omega||^2 + ||A^*\Omega||^2\}$$

for all $A \in \mathfrak{R}(\lambda\mathcal{D})$.

Furthermore, it can be shown that $c_f(\lambda)$ does not depend on the size of $\mathcal{D}$, nor the details of the theory in question. Thus for small λ, the modular automorphisms of $\mathfrak{R}(\mathcal{D})$ act approximately geometrically.

If the theory is conformally invariant, then it immediately follows that $\Delta_{\mathcal{W}}$ and $\Delta_{\mathcal{D}}$ have the same spectrum (namely, $\mathbb{R}^+$). Fredenhagen was able to prove that this equivalence holds even if the theory is only asymptotically conformal. Moreover, the spectrum of $\Delta_{\mathcal{D}}$ is independent of the choice of state since in the asymptotic regime, scale transformations act similarly to the automorphisms of an asymptotically abelian system.[12] Therefore the unitary cocycle intertwining σ_ω^t and σ_ϕ^t (with ϕ an arbitrary faithful normal state), approximately commutes with operators $A \in \mathfrak{R}(\lambda\mathcal{D})$ for sufficiently small λ. Thus the Connes invariant $S(\mathfrak{R}(\mathcal{D})) = S(\mathfrak{R}(\mathcal{W})) = [0, \infty)$, and so $\mathfrak{R}(\mathcal{D})$ must be type III_1.

[12] A C^*-system is asymptotically abelian if there exists an automorphism, α, of $\mathfrak{A}$ such that

$$||A\alpha^n(B) - \alpha^n(B)A|| \to 0 \text{ as } n \to \infty,\ \forall A, B \in \mathfrak{A}.$$

Apart from the standard Haag-Kastler axioms and the Bisognano-Wichmann relations, the crucial additional assumption needed for Fredenhagen's argument is asymptotic scale invariance. The original 1985 proof relies on the scaling properties of Wightman fields, restricting the scope of the argument to theories generated by such fields. Subsequent groundbreaking work by Buchholz and Verch (1995) on renormalization group transformations in AQFT has shown that these assumptions can be replaced by an intrinsic algebraic characterization of the asymptotic scaling properties of the observable net. The key concept they introduce is the idea of a *scaling algebra*.

Consider a net of C^*-algebras $\mathcal{O} \mapsto \mathfrak{A}(\mathcal{O})$. By rescaling the volume of regions, $\mathcal{O} \mapsto \mathfrak{A}_\lambda(\mathcal{O}) \equiv \mathfrak{A}(\lambda\mathcal{O})$ we obtain a net of C^*-algebras at scale $\lambda > 0$. The scaled net is covariant under Poincaré transformations given by $\alpha^{(\lambda)}_{\Lambda,a} = \alpha_{\Lambda,\lambda a}$. For different values of λ, the nets $(\mathfrak{A}_\lambda, \alpha_{\Lambda,\lambda a})$ are typically non-isomorphic. Thus, the scaled nets describe inequivalent theories with different energy spectra, S-matrices, charges, particle content, and coupling constants. The scaled nets will be isomorphic iff the original $\lambda = 1$ theory has full conformal symmetry.

When conformal symmetry is not present, we can still compare the structure of the theories at different scales via renormalization group transformations $R_\lambda : \mathfrak{A}(\mathcal{O}) \to \mathfrak{A}_\lambda(\mathcal{O})$. The formal properties of such transformations are characterized by a new algebraic object Buchholz and Verch call a *scaling algebra*. The scaling algebra, $\underline{\mathfrak{A}}$, associated with a theory, $\mathfrak{A}$, is an algebra of operator-valued functions $\lambda \mapsto \underline{A}_\lambda \in \mathfrak{A}$ designed to represent the orbits $\lambda \mapsto R_\lambda(A)$, $\lambda > 0$ of observables in $\mathfrak{A}$ under all possible renormalization group transformations.[13] It has the structure of a local, Poincaré covariant net of C^*-algebras on spacetime.

[13]Formally, the quasilocal scaling algebra, $\underline{\mathfrak{A}}$, is defined as the inductive C^*-limit of local scaling algebras $\underline{\mathfrak{A}}(\mathcal{O})$, which consist of functions $\lambda \mapsto \underline{A}_\lambda \in \mathfrak{A}$ subject to the following constraints:

(a) $||\underline{A}|| \equiv \sup_{\lambda>0}||\underline{A}_\lambda|| < \infty$,

(b) $\underline{A}_\lambda \in \mathfrak{A}(\lambda\mathcal{O}),\ \lambda > 0$,

(c) $||\underline{\alpha}_{\Lambda,a}(\underline{A}) - \underline{A}|| \to 0$ for $(\Lambda, a) \to (1, 0)$,

where $(\underline{\alpha}_{\Lambda,a}(\underline{A}))_\lambda \equiv \underline{\alpha}_{\Lambda,a}(\underline{A}_\lambda)$ and $\underline{\alpha}_{\Lambda,a}$ is a strongly continuous representation of the Poincaré group which acts covariantly on the net of scaling algebras.

Every state, ϕ, on $\mathfrak{A}$ has a canonical lift to a state, $\underline{\phi}$, on $\underline{\mathfrak{A}}$, and the corresponding GNS representations are unitarily equivalent.[14]

In addition to Poincaré symmetry, the scaling algebra has a group of conformal automorphisms, $\underline{\beta}_\lambda$, induced by the renormalization transformations.[15] This gives rise to a corresponding conformal action on states, $\underline{\phi}_\lambda \equiv \underline{\phi} \circ \underline{\beta}_\lambda$. In the UV limit, we consider the behavior of the net of states $(\underline{\phi}_\lambda)_{\lambda>0}$. As $\lambda \to 0$, each ultraweak limit point of the net describes a *scaling limit state* of $\underline{\phi}$. The corresponding GNS representations define *scaling limit nets*. Interestingly, each scaling limit state is a pure, Poincaré invariant vacuum state on $\underline{\mathfrak{A}}$. This allows for a simple classification of the scaling limit nets. There are three possibilities: (1) all the nets are trivial, (2) all the nets isomorphic and non-abelian, (3) the nets are non-isomorphic.

With these ideas in place, we can state the sharpened version of Fredenhagen's result obtained by Buchholz and Verch:

Theorem 4.4 (Fredenhagen, Buchholz-Verch)**.** *Assume Haag-Kastler axioms (1)-(6) and the Bisognano-Wichmann relations. If the net $\mathfrak{A}$ has a non-trivial scaling limit, then the local algebras (in the vacuum sector) are type III_1.*

Three remarks are in order. First, since the theorem depends on the geometric properties of the wedge modular objects encoded by the Bisognano-Wichmann relations, as stated it only applies to local algebras in the vacuum sector. Despite this, there is an immediate extension to any sector that is locally quasiequivalent to the vacuum, including charged sectors obeying the DHR selection criteria. It can also be extended to sectors obeying the BF selection criteria using results on sectors asymptotically connected by so-called *large translations* (Borchers and Wollenberg, 1991).

Second, there is strong physical motivation for the requirement of a non-trivial scaling limit. Case (1) corresponds to non-renormalizable theories in which the local observables are

[14]The lift is defined by $\underline{\phi}_\lambda(\underline{A}) = \phi(\underline{A}_\lambda)$, for all $\underline{A} \in \underline{\mathfrak{A}}$.

[15]For any $\mu \in \mathbb{R}^+$, $(\underline{\beta}_\mu(\underline{A}))_\lambda = \underline{A}_{\mu\lambda}$, $\lambda > 0$.

increasingly delocalized in phase space in the UV limit. Due to the uncertainty principle, we expect $||A|| > \lambda^{-1}$ for observables in $\mathfrak{R}(\lambda\mathcal{O})$, but if $||A|| \sim \lambda^{-q}$ for some $q > 1$ (rather than scaling with λ^{-1}), then there will effectively be no local observables which occupy a bounded region of phase space. As a result the local scaling nets are trivial. Such theories are expected to violate modular nuclearity and similar compactness conditions which will be discussed in more detail in the next section.

In contrast, case (2) corresponds to physically well-behaved theories whose renormalization group flow has a UV fixed point. At the fixed point, the theory behaves like a conformal field theory, justifying the scaling algebra framework as a regimentation of the scaling properties of Wightman fields. A scaling limit isomorphic to a free field net provides an algebraic characterization of *asymptotic freedom* (i.e. a Gaussian fixed-point). A scaling limit isomorphic to some other net describes the situation of *asymptotic safety* (i.e. a non-Gaussian fixed point). Case (3) is more complicated since the scaling limit varies depending on how the approach to $\lambda = 0$ is taken. Since all the scaling limit states are vacuum states, there is potentially rich physical content here (although strangely, the properties of these vacua are scale-dependent). More work in this direction is needed.

Finally, theorem 4.4 does not settle whether or not the local algebras are factors or hyperfinite. Support for the former comes from a number of physical considerations, including Wightman field constructions. If the global algebra is a factor (as it is in the vacuum sector), factorality of $\mathfrak{R}(\mathcal{D})$ can be proven under additional continuity axioms.[16] If the local algebras are hyperfinite, they can be approximated by a family of finite von Neumann algebras. This is an extremely useful property whose support comes from modular nuclearity and the split condition which we shall turn to in the next section. A classification theorem due to Haagerup (1987) establishes that there is a single hyperfinite type III_1 factor up to isomorphism. Consequently, it is widely expected that for most theories of physical interest, the internal structure of the local algebras is generic. Physical content distinguishing between

[16]See Horuzhy (1990, Thm. 1.3.6). The theorem requires strong additivity, inner/outer continuity, and primitive causality.

different QFTs comes from the relational structure of the net, not the internal algebraic structure.

4.2 Taming Entanglement

While the Reeh-Schlieder theorem and the type III property are fascinating results concerning the deep structure of QFT, they also generate a number of interpretational problems. They indicate that entanglement and non-locality are far more common in QFT than in ordinary QM. If deep down the world is so thoroughly entangled, how is it that we can talk sensibly about localized subsystems — neutrinos, hydrogen atoms, cats, galaxies — at all?

Physically, the answer lies in the fact that even though the vacuum state cannot be a product state across spacelike separated regions, in many realistic models we expect vacuum correlations to decay rapidly with distance. If the vacuum state is unique, then it is quasiequivalent to each of its subrepresentations. As a consequence, it satisfies the *cluster property*:

$$|\omega(AU(a)B) - \omega(A)\omega(U(a)B)| \to 0, \text{ as } |a| \to \infty. \tag{4.3}$$

(Here $U(a)$ represents a spacelike translation.) Thus the correlation function between any two local operators A, B tends to zero as the separation between them goes to spacelike infinity. If the theory has a mass gap, then the decay is exponential (Fredenhagen, 1985b). For separations larger than the Compton wavelength associated with the mass gap, vacuum correlations do not present a practical barrier to approximate localization in the lab. From a *conceptual* angle, however, the picture is much trickier.

By themselves the Haag-Kastler axioms are insufficient to guarantee that a QFT has well-behaved particle sectors or local thermodynamic properties. Significant work has gone into the search for new mathematically precise constraints, yielding insight into the emergent localization structure of QFT. Two of the most promising candidates, *nuclearity* and the *split condition*, already made a brief appearance in our discussion of the PCT theorem.

Both began as heuristic physical arguments, before finding surprising mathematical connections with modular theory. Their success illustrates the fruitfulness of an open-minded, cosmopolitan approach to foundational work in QFT, and provides important clues about the physical interpretation of the modular operators.

4.2.1 Nuclearity and the Split Condition

If vacuum correlations decay exponentially, then allowing some finite buffer zone between the system and the environment should enable us to isolate the system in the region up to good approximation. Ideally, the system and environment algebras should be *statistically independent*. Two algebras are statistically independent if given any two states on $\mathfrak{R}(\mathcal{O}_1)$, $\mathfrak{R}(\mathcal{O}_2)$, there exists a common extension of these local states to a product state on $\mathfrak{R}(\mathcal{O}_1) \vee \mathfrak{R}(\mathcal{O}_2)$. (We also require that if the original states are faithful, the extension is faithful too.) Let $\mathcal{O} \subset \tilde{\mathcal{O}}$ be regions of spacetime. Statistical independence between $\mathfrak{R}(\mathcal{O})$ and $\mathfrak{R}(\tilde{\mathcal{O}})'$ is ensured if there exists an interpolating type I factor $\mathfrak{R}(\mathcal{O}) \subset \mathfrak{N} \subset \mathfrak{R}(\tilde{\mathcal{O}})'$. The *split condition* is said to hold if such a factor exists for the local algebras associated with all regions $\mathcal{O} \subset \tilde{\mathcal{O}}$.

In fact, since there exists a common cyclic separating vector for $\mathfrak{R}(\mathcal{O})$, $\mathfrak{R}(\tilde{\mathcal{O}})$, and $\mathfrak{R}(\mathcal{O})' \wedge \mathfrak{R}(\tilde{\mathcal{O}})$ (by the Reeh-Schlieder theorem), whenever there is an interpolating type I factor between $\mathfrak{R}(\mathcal{O})$ and $\mathfrak{R}(\tilde{\mathcal{O}})$ we have what is called a *standard split inclusion*. In this case, the following result applies:

Theorem 4.5. *Let $\mathfrak{M} \subset \tilde{\mathfrak{M}}$ be factors acting on $\mathcal{H}$ such that there exists a common cyclic, separating vector, Ω, for $\mathfrak{M}$, $\tilde{\mathfrak{M}}$, and $\mathfrak{M}' \wedge \tilde{\mathfrak{M}}$. Then the following conditions are equivalent:*[17]

(i) $\mathfrak{M}$ and $\tilde{\mathfrak{M}}'$ are statistically independent,

(ii) there exists a vector, $\eta \in \mathcal{H}$, cyclic and separating for $\mathfrak{M} \wedge \tilde{\mathfrak{M}}'$ such that

$$\langle \eta | AB' | \eta \rangle = \langle \Omega | A | \Omega \rangle \langle \Omega | B' | \Omega \rangle, \;\; A \in \mathfrak{M}, \;\; B' \in \tilde{\mathfrak{M}}',$$

[17] See Haag (1996, Prop. 5.2.1) for a concise proof.

(iii) there exists a unitary operator $W : \mathcal{H} \to \mathcal{H} \otimes \mathcal{H}$ *such that,*

$$WAB'W^* = A \otimes B',\ A \in \mathfrak{M},\ B' \in \tilde{\mathfrak{M}}',$$

(iv) there exists a type I factor, $\mathfrak{N}$*, such that* $\mathfrak{M} \subset \mathfrak{N} \subset \tilde{\mathfrak{M}}$*.*

Furthermore, the modular objects of $\mathfrak{M}$ and $\tilde{\mathfrak{M}}$ are interrelated and can be used to naturally define $\mathfrak{N}$:

Theorem 4.6 (Doplicher-Longo)**.** *If* $\Lambda = \{\mathfrak{M}, \tilde{\mathfrak{M}}, \Omega\}$ *is a standard split inclusion, then there exists a canonical choice of interpolating type I factor* $\mathfrak{M} \subset \mathfrak{N}_\Lambda \subset \tilde{\mathfrak{M}}$*.*[18]

Thus if $\varphi : \Lambda_1 \to \Lambda_2$ is an isomorphism of standard split inclusions, then $\varphi(\mathfrak{N}_{\Lambda_1}) = \mathfrak{N}_{\Lambda_2}$. Since $\mathfrak{M}$ and $\tilde{\mathfrak{M}}$ are factors (though only one needs to be), $\mathfrak{N}_\Lambda$ has an elegant characterization in terms of the modular conjugation, J, associated with the algebra $\mathfrak{M}' \wedge \tilde{\mathfrak{M}}$:

$$\mathfrak{N}_\Lambda = \mathfrak{M} \vee J\mathfrak{M}J = \tilde{\mathfrak{M}} \wedge J\tilde{\mathfrak{M}}J \tag{4.4}$$

Moving back to the spacetime picture, let $\mathcal{D} \subset \tilde{\mathcal{D}}$ be an inclusion of double cones. Assuming duality, we have $\mathfrak{R}(\mathcal{D})' = \mathfrak{R}(\mathcal{D}')$ and $\mathfrak{R}(\tilde{\mathcal{D}})' = \mathfrak{R}(\tilde{\mathcal{D}}')$. The algebra $\mathfrak{R}(\mathcal{D})' \wedge \mathfrak{R}(\tilde{\mathcal{D}}) \equiv \mathfrak{R}(\mathcal{L})$ is the local algebra associated with the shaded region (denoted $\mathcal{L}$). We'll refer to it as the *life-preserver* region because it forms a protective ring around $\mathcal{D}$, allowing us to isolate $\mathfrak{R}(\mathcal{D})$ from the environment represented by $\mathfrak{R}(\tilde{\mathcal{D}}')$. Using relation (4.4), we can view $\mathfrak{N}_\Lambda$ as the algebra generated by $\mathfrak{R}(\mathcal{D})$ and $J\mathfrak{R}(\mathcal{D})J$. As we'll go on to see in the next subsection, the structure of $\mathfrak{N}_\Lambda$ is crucial for defining local analogues of certain global quantities including charge observables and the Hamiltonian.

The split property ensures that we can find product states over $\mathfrak{R}(\mathcal{O})$ and $\mathfrak{R}(\tilde{\mathcal{O}}')$ given any local state in $\mathcal{O}$. It is therefore possible in principle to disentangle a system localized in $\mathcal{O}$ from the distant environment (i.e. the spacelike complement of the enveloping region

[18] Doplicher and Longo (1984, Thm. 2.1)

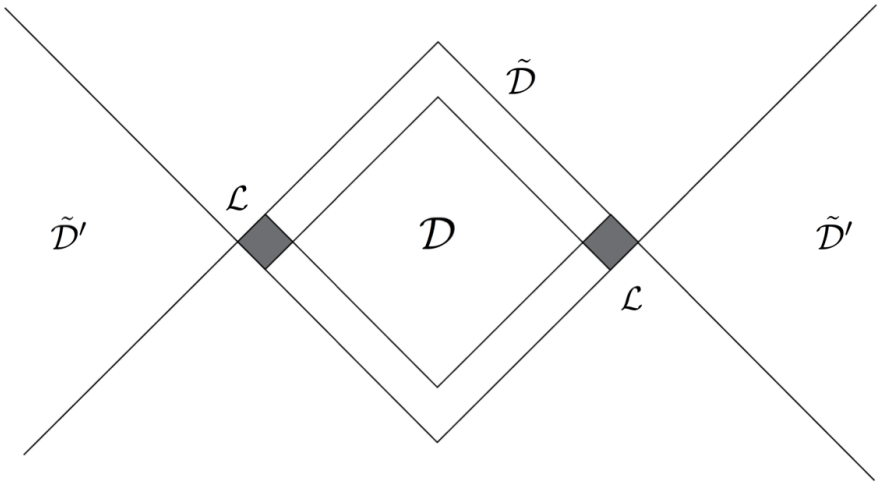


Figure 5: A split inclusion $\Lambda = \{\mathfrak{R}(\mathcal{D}), \mathfrak{R}(\tilde{\mathcal{D}}), \Omega\}$

$\tilde{\mathcal{O}}$).[19] In addition, since any type I factor is hyperfinite, the split property entails that the local algebras are also hyperfinite, and thus can be approximated by an ascending sequence of finite algebras.

While these features are extremely useful, they are not enough to guarantee the existence of local KMS states and other desirable thermodynamic properties. In classical field theory, states which are bounded in space and energy occupy a finite-volume region of phase-espace. By analogy, in QFT one might hope to find conditions that ensure states which are Knight-localized with bounded energy content will occupy a finite-dimensional subspace of $\mathcal{H}$. *Nuclearity*, originally devised by Buchholz and Wichmann (1986), is a condition designed to secure this result.[20]

Recall that Knight-localized states can be generated by applying local norm-1 operators to the vacuum. Given a double-cone of radius r, we let $\mathfrak{R}^{(1)}(\mathcal{D})$ denote the unit ball in $\mathfrak{R}(\mathcal{D})$. Additionally, we impose an exponential energy damping term $e^{-\beta H}$ to restrict the energy (here H is the global Hamiltonian and $\beta > 1$ is an inverse temperature constant). Nuclearity

[19] As the life-preserver region shrinks, however, the amount of energy required to disentangle the system grows.

[20] Nuclearity was originally inspired by the Haag-Swieca compactness condition (Haag and Swieca, 1965).

amounts to the requirement that

$$\mathcal{N}_{\beta,r} = e^{-\beta H}\mathfrak{R}^{(1)}(\mathcal{D})\Omega, \tag{4.5}$$

is a nuclear set in $\mathcal{H}$. This means that there is a positive trace class operator $T_{\beta,r}$ such that $\mathcal{N}_{\beta,r} = T_{\beta,r}\mathcal{H}^{(1)}$, where $\mathcal{H}^{(1)}$ is the unit ball in $\mathcal{H}$. This ensures that the subspace $\mathcal{N}_{\beta,r}$ is "essentially finite dimensional." There is an ascending sequence of finite dimensional subspaces $\mathcal{H}_{d(\epsilon)} \subset \mathcal{H}$, with dimension $d(\epsilon)$, $\epsilon > 0$, such that vectors in $\mathcal{N}_{\beta,r}$ which are orthogonal to $\mathcal{H}_{d(\epsilon)}$ have length less than ϵ. It is important to emphasize that $\mathcal{N}_{\beta,r}$ is not a closed subspace of $\mathcal{H}$, in fact $\overline{\mathcal{N}_{\beta,r}} = \mathcal{H}$. Rather, the idea is that the length of vectors in $\mathcal{N}_{\beta,r}$ grow with increasing delocalization, quantified by the ratio $||A||/||A\Omega||$. This permits us to estimate phase space volumes corresponding to localized states with bounded energy.

The dependence of $d(\epsilon)$ on β and r is characterized by the *nuclearity index*, $\nu_{\beta,r} \equiv \inf$ (tr $T_{\beta,r}$) (where the infimum is taken over all trace class operators such that $\mathcal{N}_{\beta,r} = T_{\beta,r}\mathcal{H}^{(1)}$). If the nuclearity index satisfies the bound $\nu_{\beta,r} < \exp cr^3\beta^{-n}$ (for positive constants c, n) then nuclearity entails the split condition for all sets of nested double-cones. If the nuclearity index diverges, then the split condition only holds if the separation between $\mathcal{O}$ and $\tilde{\mathcal{O}}$ is larger than the inverse Hagedorn temperature. In the Standard Model, this temperature refers to the point at which hadronic matter is no longer stable since there is sufficient energy for creating quark-antiquark pairs from the vacuum. This scale is typically thought to be at least two orders of magnitude smaller than the Planck scale, hence we can expect the split property to obtain for all systems within current experimental accessibility even if the nuclearity bound is not met.

As formalized by Buchholz and Wichmann, the nuclearity condition is defined using the global Hamiltonian. A purely local version of nuclearity was obtained by Buchholz et al. (1990), by exploiting properties of the local modular operators. *Modular nuclearity* requires

the map

$$\Xi A \equiv (1 + \Delta^{-1/2})^{-1} A\Omega, \;\; A \in \mathfrak{R}(\mathcal{D}) \tag{4.6}$$

to be nuclear. This provides a bound on the same order of magnitude as the original nuclearity condition. In addition to being a potentially useful generalization of nuclearity applicable in both flat and curved spacetime settings, modular nuclearity reveals a deep connection between the local energy and the modular operator Δ. Two observers estimating the phase space volume of localized, energy-bounded states using H and Δ will agree. This has led a number of commentators to speculate that the *modular Hamiltonian*, $\ln\Delta$, which acts as the generator of the modular automorphism group, represents a kind of local dynamics for the region $\mathcal{D}$. This question dovetails with the Connes-Rovelli *thermal time hypothesis* and will be taken up at the end of Ch. 5.

4.2.2 Local Charge Structure

Although the localization of charge is built into the DHR and BF analysis of superselection structure, neither provides a means of constructing local observables measuring the net charge in a spatiotemporal region. Sectors are labeled by the eigenvalues of global charge observables — unbounded operators affiliated with the center of the global algebra, $\pi(\mathfrak{A})''$, a type I von Neumann algebra. But the local algebras are type III factors, so their center is trivial. It is not at all obvious how to find local analogues of the global charges. A similar problem arises in the case of the Hamiltonian. The dynamics in AQFT are specified by a strongly continuous unitary representation of the translation subgroup acting on the global Hilbert space. Stone's theorem ensures that there are self-adjoint operators in $B(\mathcal{H})$ acting as generators of the translations. Fixing a global cauchy surface then suffices to single out a generator of the time translations, the global Hamiltonian. By themselves, the Haag-Kastler axioms do not guarantee that there is a well-defined local Hamiltonian; a locally measurable observable implementing the dynamics in a given subregion of spacetime. Of course we

routinely measure localized charges and energies in laboratory experiments. How is it then, that these globally defined quantities can be connected to the local observable structure of AQFT?

In a series of creative papers, Doplicher, Longo, and Buchholz developed a technique for constructing local charge operators utilizing the split property.[21] The construction employs the unitary operator $W : \mathcal{H} \to \mathcal{H}\otimes\mathcal{H}$ whose existence is guaranteed by theorem 4.5.iii. Fixing a target region $\mathcal{O}$ and an enveloping region $\tilde{\mathcal{O}}$, by theorem 4.6 there exists a canonical type I interpolating factor $\mathfrak{R}(\mathcal{O}) \subset \mathfrak{N}_\Lambda \subset \mathfrak{R}(\tilde{\mathcal{O}})$. The proof of theorem 4.5 shows that given any interpolating type I factor, W can be chosen so that the image of $B(\mathcal{H}) \otimes I$ is that factor. Thus in the present case, there exists a unitary W such that $W^*(B(\mathcal{H}) \otimes I)W = \mathfrak{N}_\Lambda$. This defines a *universal localization map* associated with the standard split inclusion Λ:

$$\psi_\Lambda(A) = W^*(A \otimes I)W \in \mathfrak{N}_\Lambda,\ \forall A \in B(\mathcal{H}). \tag{4.7}$$

In the DHR/BF framework, $\mathcal{H}$ will be the large Hilbert space on which the field algebra, $\mathfrak{F}$, acts irreducibly. Given a compact gauge group, G, there is a unitary representation, $U(g) \in B(\mathcal{H})$, of G acting on $\mathcal{H}$ such that,

$$U(g)\mathfrak{F}(\mathcal{O})U(g)^{-1} = \alpha_g\mathfrak{F}(\mathcal{O}) = \mathfrak{F}(\mathcal{O}). \tag{4.8}$$

The superselection sectors correspond to the invariant subspaces of $\mathcal{H}$ under the action of $U(g)$. The global algebra of observables for each sector can either be viewed as the fixed points of $\mathfrak{F}$ under the gauge action, or equivalently, as unitarily inequivalent, irreducible representations of the quasilocal algebra, $\mathfrak{A}$, meeting the DHR/BF selection criteria.

[21] Doplicher (1982) proved an early version of the result for abelian gauge charges only. Subsequently, Doplicher and Longo (1983) extended the analysis to cover non-abelian gauge symmetries. Buchholz et al. (1986) introduced the concept of the universal localization map and provided a generalization to spacetime symmetries and supersymmetries.

In this context, given a nested set of spatiotemporal regions, the universal localization map finds a structural analogue of $B(\mathcal{H})$ in terms of operators localized in the outer region, $\tilde{\mathcal{O}}$. In particular, $\psi_\Lambda(U(g)) = U_\Lambda(g)$ yields a strongly continuous local representation of the gauge group with the following features:

$$\begin{aligned}
&\text{(i)} \quad U_\Lambda(g) \in \mathfrak{F}(\tilde{\mathcal{O}}) \\
&\text{(ii)} \quad U_\Lambda(g) F U_\Lambda(g)^{-1} = \alpha_g(F),\ F \in \mathfrak{F}(\mathcal{O}) \\
&\text{(iii)} \quad \alpha_h(U_\Lambda(g)) = U(hgh^{-1}),\ h, g \in G
\end{aligned}$$

This enables us to find unitary operators in the outer field algebra, $\mathfrak{F}(\tilde{\mathcal{O}})$, which implement gauge transformations on the inner field algebra $\mathfrak{F}(\mathcal{O})$. If J is a self-adjoint operator in $\mathfrak{F}$ which acts as one of the generators of $U(g)$, then $\psi_\Lambda(J) = J_\Lambda \in \mathfrak{F}(\tilde{\mathcal{O}})$ is a local generator for $U_\Lambda(g)$. Moreover if Q is a global charge observable contained in $Z(\pi(\mathfrak{A})'')$, i.e. a Casimir invariant of $U(g)$, then $\psi_\Lambda Q = Q_\Lambda$ is a local charge observable measuring the charge in $\mathcal{O}$, a Casimir invariant of $U_\Lambda(g)$. In general Q_Λ will be an unbounded observable affiliated with $\mathfrak{R}(\mathcal{L})$, and therefore has a spectral decomposition consisting of self-adjoint projections in $\mathfrak{R}(\mathcal{L})$. (The local generators J_Λ will themselves be observables iff the gauge group is abelian.) We will call observables affiliated with the life-preserver region around $\mathcal{O}$, *semilocal* observables

Nothing in this construction relies on the group representation being an internal gauge symmetry. Given a global representation of the translations, the universal localization map defines a corresponding local representation, $U_\Lambda(a)$, in terms of semilocal unitaries. Since the translations commute with the gauge group, the generators of $U_\Lambda(a)$ are observables in $\mathfrak{R}(\mathcal{L})$. These include a semilocal Hamiltonian operator $\psi_\Lambda H = H_\Lambda$. The only notable difference from the gauge case is the fact that the semilocal Hamiltonian only generates time translations small enough to preserve localization within the inner region, $\mathcal{O}$.

Thus with the help of the split condition, one can identify semilocal generators of global symmetries which are measurable in a region slightly larger than the one they act on. While these initial results are encouraging, more mathematical and interpretive work is needed to understand the features of these observables. One oddity is the fact that the semilocal observables are unbounded, yielding a true representation of the global group. From examples in non-relativistic QM as well as Wightman field theory, we would expect the local Hamiltonian and local charge operators for a compact region to be bounded, viewed as an integral of a local energy/charge density over the target region, $\mathcal{O}$. As a result H_Λ and Q_Λ cannot be interpreted similarly. Another related issue concerns the convergence properties of the semilocal observables. D'Antoni et al. (1987) have shown that under the assumption of modular nuclearity, the semilocal observables converge (strongly) to their global counterparts as $\tilde{\mathcal{O}}$ expands to cover all of spacetime. The short distance limit is more problematic. The definition of the universal localization map depends crucially on the choice of the enveloping region $\tilde{\mathcal{O}}$. One would like to show that as $\tilde{\mathcal{O}} \to \mathcal{O}$, the semilocal observables converge, but at this stage we have poor control of the limit. The collection of canonical type I factors, $\{\mathfrak{N}_\Lambda\}$, does have a well-defined net structure, however, making it a potentially rich object for further study.[22]

The relation $\mathfrak{N}_\Lambda = \mathfrak{R}(\mathcal{D}) \vee J\mathfrak{R}(\mathcal{D})J$ seems to point to the following physical interpretation of the construction detailed above: in order to define local charge observables or a local Hamiltonian for a double-cone region, we need to take into account correlations between local observables in $\mathfrak{R}(\mathcal{D})$ and the nearby environment. By duality we have $\mathfrak{R}(\mathcal{L}) = \mathfrak{R}(\mathcal{D} \wedge \tilde{\mathcal{D}}')$, thus $J\mathfrak{R}(\mathcal{D})J \subset \mathfrak{R}(\mathcal{L})$ which we can view as the extra observables in the life-preserver region needed to regularize the charge operators and Hamiltonian for $\mathcal{D}$.

This conclusion has interesting interpretive ramifications for the concept of a "local observable." Traditionally the elements of $\mathfrak{R}(\mathcal{O})$ are viewed as the physical quantities which

[22]Let Λ_1, Λ_2 be any two standard split inclusions of local algebras (with a common cyclic, separating vector Ω). If $\tilde{\mathcal{O}}_1 \subset \tilde{\mathcal{O}}_2$, then $\mathfrak{N}_{\Lambda_1} \subset \mathfrak{N}_{\Lambda_2}$, providing an analogue of isotony. If $\tilde{\mathcal{O}}_1 \subset \tilde{\mathcal{O}}_2'$, then $\mathfrak{N}_{\Lambda_1} \subset \mathfrak{N}_{\Lambda_2}'$, yielding a version of microcausality. The net is also Poincaré covariant and fulfills analogues of duality and additivity. See Doplicher and Longo (1983, Thm. 4.1) for further details.

can be measured by experiments causally confined to $\mathcal{O}$. A number of philosophers (Wallace, 2012; Ruetsche, 2011a) have adopted a less operationalist interpretation — the elements of $\mathfrak{R}(\mathcal{O})$ represent the quantities which characterize the local properties quantum field in region $\mathcal{O}$. But the present results show that many of the most important quantities characterizing the physics inside $\mathcal{O}$ are not in fact localized in $\mathcal{O}$ but rather the life-preserver region around it. This puts some pressure on the latter reading. The realist could retreat and say that $\mathfrak{R}(\mathcal{O})$ only contains some of the quantities describing the local physics, or press the attack and shift attention to $\mathfrak{N}_\Lambda$. Of course, $\mathfrak{N}_\Lambda$ isn't a local algebra at all! It is impossible to find a region of spacetime whose local operators include all and only elements of $\mathfrak{N}_\Lambda$. (Of course according to the operationalist story, they are measurable in the vicinity of $\mathcal{O}$). Either way, the realist's hand is forced. She must provide a more nuanced story about the localization of observables.

Another pressing issue is how the existence of local charge observables squares with various no-go theorems militating against position observables in QFT. Malament (1996) proves that there cannot be a translation covariant, projection-valued measure describing sharp spatial localization in relativistic QM. Such a measure would correspond to the spectral decomposition of a self-adjoint position observable. Halvorson and Clifton (2002) extend this result to positive operator-valued measures, ruling out the possibility of a symmetric (non-self-adjoint) position operator for unsharply localized particles. These theorems are important pieces of the broader argument against particle interpretations of QFT. If a particle were ever (sharply or unsharply) localized in some region, $\mathcal{S}$, of a spatial hypersurface at time t, then instantaneously thereafter, the particle would have some probability of being detected anywhere else in the universe. To avoid direct conflict with relativistic causality, the conclusion drawn is that nothing can be so localized in the first place. But if this is true, how is it possible for a *charge* to be localized, even approximately, and how can we measure it with semilocal observables?

Fortunately for experimental physicists everywhere, the conflict is only apparent.[23] The Malament-Clifton-Halvorson no-go results do not apply to local charge observables. The reason lies in the central assumption Malament calls *localizability*. If $\mathcal{S}_1$, $\mathcal{S}_2$ are spatial regions on the same hypersurface separated by finite distance and E_1, E_2 are the corresponding localization projections/effects, then if $\phi(E_1) = 1$, localizability demands that $\phi(E_2) = 0$. For projection operators, this condition is equivalent to orthogonality $E_1E_2 = 0$. For the more general POVM case considered by Halvorson and Clifton, localizability is equivalent to $E_1 + E_2 \leq I$.[24] Localizability has a direct 4-d analogue: for any two spatiotemporal regions, $\mathcal{O}_1$, $\mathcal{O}_2$, spacelike separated by a finite distance, $\phi(E_1) = 1 \Rightarrow \phi(E_2) = 0$. The motivation is clear — the probability of detecting the same localized particle in two spacelike separated regions cannot be greater than one if the particle cannot travel faster than light.

But there is no reason why local charge observables should obey Malament's localizability condition. Let Q_1, Q_2 be local charge observables for spacelike separated regions $\mathcal{O}_1$,$\mathcal{O}_2$, and $\{E_{Q1}\}$, $\{E_{Q2}\}$ their spectral decomposition (the projections will be self-adjoint operators in $\mathfrak{R}(\mathcal{L}_1)$, $\mathfrak{R}(\mathcal{L}_2)$ respectively). If $E_{Q1}(+4)$ is the projection taking eigenvalue 1 if the charge in $\mathcal{O}_1$ is +4 and zero otherwise, and $E_{Q2}(+4)$ is the corresponding projection for $\mathcal{O}_2$, there is no reason why these projections should be orthogonal (or if we generalize to effects, why $E_{Q1}(+4) + E_{Q1}(+4) \leq I$). The probabilities of measuring some charge in $\mathcal{O}_1$ and the same charge in $\mathcal{O}_2$ are completely independent of one another.[25]

The position observables targeted by the no-go theorems and the charge observables we've been discussing here are very different kinds of objects. The former describe the position of a *particular* localized particle. The latter describe the net charge in a region of spacetime.

[23] It is tempting to think that the issue lies in the difference between spatial and spatiotemporal localization, but this turns out to be a red herring. There are parallel no-go theorems for spatiotemporal position observables. See Giannitrapani (1998).

[24] In order to prove their theorem, Halvorson and Clifton actually impose a slightly stronger condition, *additivity* which requires that $E_1 + E_2 = E_{12}$ where E_{12} is the localization effect for the region $\mathcal{S}_1 \cup \mathcal{S}_2$.

[25] Once we drop localizability, the theorem fails to get off the ground. In Malament's proof, it's the fact that the projection E_1E_2 vanishes initially that allows us to apply Hegerfeldt's lemma — since the Hamiltonian is half-bounded either E_1E_2 vanishes for all t or it's non-zero for almost all t — to arrive at a contradiction. (Halvorson and Clifton's proof uses an argument by induction where the base case is shown to satisfy the conditions of Malament's theorem.)

The same particle cannot be two places at once, but two different regions can easily have the same net charge.

But if the no-go theorems are consistent with localized charges, then what about the broader conclusion drawn from these theorems — there cannot be sharply or unsharply localized entities of any sort in QFT? Here is it crucial to remember the original question that Malament's paper seeks to answer: "Can there be a relativistic quantum theory describing a *single* sharply localized particle?" The extension to a finite number of particles is obvious, assuming that we can write the Hilbert space for the multi-particle system as a tensor product of one-particle Hilbert spaces. Halvorson and Clifton's paper covers both sharply and unsharply localized particles. Together these no-go results point towards a negative answer to the question, vindicating the dogma that Malament seeks to defend: the unification of QM and relativity requires fields.

But QFT adopts this dogma from the start; it is the theory of relativistic quantum *fields* after all. Even if one balks at the idea of a fundamental field ontology, at a minimum QFT requires the possibility of an infinite number of particles with spontaneous pair creation/annihilation from the vacuum. But the no-go theorems are effectively silent about the possibility of localized entities in theories of this kind. The reason is that the theorems assume that we can index observables by the individual particles they describe. For each localized particle there will be a position observable *for that particle*. In a theory describing an infinite number of particles, however, we cannot consistently do this. If the single particle Hilbert spaces have dim>1, then taking an infinite tensor product will generate a non-separable global Hilbert space. In all of its extant variants, QFT assumes that the global Hilbert space is separable, so without modifying the formalism, there is no way to consistently index observables by particle. Even if we move to a non-separable Hilbert space, it's not clear that such indexing would be possible. If a pair of particles annihilate in region and another pair create in a spacelike separated region, which particles are they?

It may be possible to construct a Malament-style no-go argument against an infinite particle theory framed in a non-separable Hilbert space. For present purposes, the important point is that even if we adopt Malament's dogma, the no-go theorems do not rule out the possibility of approximately localized field quanta, despite appearances to the contrary.[26] Taken at face value, the mathematical formalism of QFT doesn't describe the interaction of multiple individual systems, rather it tells us about the distribution of properties like mass, spin, charge, and energy across spacetime. With the help of nuclearity and the split condition, it is possible to approximately localize some of these properties within particular spatiotemporal regions. We can think of stable bundles of these localized properties as describing emergent particles — quanta of an underlying fundamental field (the details of which the theory is largely silent about).

Despite our lack of a traditional position observable, nothing prevents us from tracking the motion of stable particles throughout spacetime. For example, a sequence of "yes" measurements of local charge $Q_\Lambda = -1$ in timelike related regions gives us a way to track the motion of an electron. This basic idea underwrites our use of a cloud chambers and scintillator detectors for measuring the trajectories of charged particles. Of course any local positive operators that we use to track the location of a particle in QFT must have non-vanishing expectation in the vacuum as a consequence of the Reeh-Schlieder theorem. So in our simple electron tracking example, it's possible that when one of our scintillators fires, indicating the presence of -1 charge in a certain part of the detector, another scintillator also fires in some spacelike separated region of the detector. But the reason for this isn't that particles can travel at superluminal velocities. Rather it's that any local observable can accidentally respond to vacuum excitations. Through careful engineering we can lower the probability of a misfire, but we can never make it vanish entirely.[27]

[26] No-go theorems for local number operators might be an exception. See for example, Halvorson and Clifton (2002, Thm. 3).

[27] These considerations provide a partial response to Barrett (2002). We *can* find localized measurement records in QFT. The localized charge operators Q_Λ are a prime example. One would like to show that even though there is no position observable, it is possible to construct a relative position observable measuring

The dogma (at least in its weak form) is safe. There is no relativistic quantum theory of finitely-many localized particles. But as with all dogmas, we must be wary of far reaching conclusions drawn from it.

4.3 Spacetime State Realism

In recent work, David Wallace and Christopher Timpson present a new ontological picture of quantum theory dubbed *spacetime state realism* (Wallace and Timpson, 2010; Wallace, 2012). In contrast to both wavefunction realism and primitive ontology views, spacetime state realism posits a field of local density operators as its fundamental object. The basic idea is to assume that the universe can be broken up into subsystems, $\{S_i\}$, represented by a tensor product decomposition of the global Hilbert space:

$$\mathcal{H} = \bigotimes_{S_i} \mathcal{H}_{S_i} \tag{4.9}$$

The density operator associated with a given subsystem, S_a, is then defined by the partial trace of the global state over all $\mathcal{H}_{S_i}$, $i \neq a$. Wallace and Timpson propose that this density operator represents the subsystem's intrinsic properties, "just as the field values assigned to each spacetime point in electromagnetism represented the (electromagnetic) intrinsic properties which that point instantiated" (Wallace, 2012, p. 299).

The picture just sketched describes a collection of density operators indexed by subsystem. Strictly speaking, nothing requires that subsystems be individuated by their spatiotemporal properties, but the view takes on a more concrete, field-theoretic character under this assumption. Taking their cue from AQFT, Wallace and Timpson consider the net of local observable algebras $\mathfrak{A}(\mathcal{O})$, assigned to bounded spacetime regions $\mathcal{O}$. Given a local state, ρ (viewed as a linear functional on $\mathfrak{A}(\mathcal{O})$), one can construct a local Hilbert space $\mathcal{H}_\rho$ via

the distance between two localized charges. In principle this could be achieved through a description of measuring rods as composite systems composed of stable localized charges.

the GNS construction. ρ is then represented by a density operator (a vector in fact) on $\mathcal{H}_\rho$. According to Wallace and Timpson, the basic ontology of AQFT is thus a field of local density operators indexed by spacetime region. As they go on to explain, this viewpoint has a number of virtues:

> This alternative ontology avoids the problems previously identified for wave-function realism: it is well defined for any quantum theory with compositional structure (in particular, for any quantum field theory, and for any many-particle theory once it is expressed in field-theoretic terms); it respects the dynamical structure of QM, indicating no preference for Schrödinger over Heisenberg or interaction dynamics (as the state is just construed as a linear functional of the dynamical variables); it adds no additional interpretational structure (given that the compositional structure of the system is, *ex hypothesi*, already contained within the formalism); and it gives an appropriately central role to spacetime. For want of a better name, we call it *spacetime state realism*. (Wallace and Timpson, 2010, p. 712)

The results concerning the type of local algebras discussed in §4.1 pose a significant technical challenge to this picture. The core of spacetime state realism is the idea that the intrinsic properties of spatiotemporally individuated subsystems can be represented by density operators acting on separate subspaces of a global Hilbert space. As we'll go on to see, however, this viewpoint tacitly relies on intuitions about type I algebras which are invalid in AQFT.

First off, there is something rather odd about Wallace and Timpson's proposal for AQFT. The fundamental ontology they propose consists of a field of local density operators assigned to each region of spacetime, but in practice, these density operators are defined in terms of other objects, the net of local C^*-algebras and the global state, ϕ, (viewed as a linear functional over the quasilocal algebra $\mathfrak{A}$). By taking the restriction of ϕ to a particular local algebra, $\mathfrak{A}(\mathcal{O})$, we obtain a local state, ρ, and thus a local density operator and Hilbert

space via the GNS construction. Note that due to the non-separability of the quantum state, specifying the local states over an arbitrary partition of spacetime is insufficient to determine the global state. So prima facie, the local state has derivative rather than fundamental status in AQFT.

As Baker (2014) notes, Wallace's comments elsewhere help shed light on this puzzling situation. In the opening chapter of *The Emergent Multiverse*, Wallace writes

> In the so-called 'algebraic' description of QFT, an algebra of operators is associated to each spatial region, so that the operators associated with region $\mathcal{O}$ are intended to represent observables localized in $\mathcal{O}$. At least formally, we can regard this as equivalent to decomposing the Hilbert space into quotient spaces, each one representing the quantum state of a different spatial region. (Wallace, 2012, p.15)

Wallace goes on to reference one of his older papers (Wallace, 2006), where he fleshes out an argument for the formal equivalence of these two procedures in more detail. First, partition spacetime into disjoint subsets $\{\mathcal{O}\}$. Then for each element of the partition, define the Hilbert space $\mathcal{H}_\mathcal{O}$ as the space of complex functionals on functions on $\mathcal{O}$. Since any function on spacetime can be specified uniquely by its restriction to each region in $\{\mathcal{O}\}$, we can write $\mathcal{H} = \bigotimes \mathcal{H}_{\mathcal{O}_i}$. Given a collection of local field operators, $\{\psi(f)\}$, where f has compact support in $\mathcal{O}$, it follows that $\{\psi(f)\}$ act trivially outside of $\mathcal{O}$, thus we can view $\mathcal{H}_\mathcal{O}$ as the local Hilbert space representing subsystems of the field localized in $\mathcal{O}$. Alternatively, Wallace notes that we can start with the subspace $\mathcal{H}_\mathcal{O}$ and define the local field operators as those (unbounded) operators on $\mathcal{H}$ which act trivially everywhere except $\mathcal{O}$.

Thus it is clear that Wallace and Timpson take the standard observable-first procedure for defining the local density operators as a pragmatic convention. In principle we could run things backwards, starting with local states/Hilbert spaces. Moreover, it is this order of definition that reflects the order of grounding according to spacetime state realism.

This is where things go awry once we pay closer attention to the mathematical details. To start, if we treat each bounded, spacelike separated region as a separate subsystem, we immediately run into a familiar problem. Assuming the local Hilbert space for each region has dim>1, then the infinite tensor product $\bigotimes \mathcal{H}_{\mathcal{D}}$ is a nonseparable Hilbert space, contradicting explicit assumptions both in AQFT and Wightman QFT. Wallace and Timpson's proposal only works when the division of spacetime into subsystems is a finite partition. In the simplest case, we can consider a partition consisting of a local subsystem region $\mathcal{D}$ and its spacelike environment $\mathcal{D}'$. As we have seen in §4.1, the local algebras $\mathfrak{R}(\mathcal{D})$ and $\mathfrak{R}(\mathcal{D}') = \mathfrak{R}(\mathcal{D})'$ are type III_1 factors.[28] From the observable-first standpoint, dividing the universe into spatiotemporal subsystems amounts to breaking up $B(\mathcal{H})$ into type III_1 subalgebras. But this procedure is not at all formally equivalent to splitting $\mathcal{H}$ into a tensor product:

Proposition 4.7. *If $\mathfrak{M} \subset B(\mathcal{H})$ is a type III factor, then $\mathfrak{M}$ and $\mathfrak{M}'$ cannot be split, i.e. there does not exist an isomorphism $V : \mathcal{H} \to \mathcal{K}_1 \otimes \mathcal{K}_2$ such that $V^*\mathfrak{M}V \subset B(\mathcal{K}_1) \otimes I$ with $V^*\mathfrak{M}'V \subset I \otimes B(\mathcal{K}_2)$.*[29]

Thus we cannot split $\mathcal{H}$ into a tensor product such that the local algebra $\mathfrak{R}(\mathcal{D})$ acts on one part and $\mathfrak{R}(\mathcal{D}')$ acts on the other. Proposition 4.7 tells us that peculiarities of infinite tensor products aside, we still can't decompose the global Hilbert space into a local system portion plus an environment portion as we can in non-relativistic QM.[30]

Perhaps requiring a tensor product decomposition of $\mathcal{H}$ is too strong. In the operator-first proposal, we took the restriction of the global state to $\mathfrak{A}(\mathcal{O})$. Then via the GNS construction, we obtained a local density operator, D_ρ, acting on the local GNS Hilbert space, $\mathcal{H}_\rho$. Perhaps the spacetime state realist can exploit these local Hilbert spaces to ground his ontology even

[28] The global algebra $\pi(\mathfrak{A})''$ must be type I_∞ due to the spectrum condition and microcausality. If the theory has gauge symmetry, it will not be a factor except in the vacuum sector. See Araki (1964a).

[29] *Proof.* Assume such an isomorphism exists. Since $V^*\mathfrak{M}V \subset B(\mathcal{K}_1) \otimes I$, it follows that $(B(\mathcal{K}_1) \otimes I)' \subset (V^*\mathfrak{M}V)'$. But $(B(\mathcal{K}_1) \otimes I)' = I \otimes B(\mathcal{K}_2)$ and $(V^*\mathfrak{M}V)' = V^*\mathfrak{M}'V$, so $I \otimes B(\mathcal{K}_2) \subset V^*\mathfrak{M}'V$. But by hypothesis $V^*\mathfrak{M}'V \subset I \otimes B(\mathcal{K}_2)$, so $V^*\mathfrak{M}'V = I \otimes B(\mathcal{K}_2)$. But this is impossible since $V^*\mathfrak{M}'V$ is type III and $I \otimes B(\mathcal{K}_2)$ is type I. □

[30] The split condition can certainly help isolate subsystems in QFT, but it won't help Wallace. See §4.3.3.

if $\mathcal{H}$ cannot be decomposed into a tensor product. Alas, this idea is short-lived. As it turns out, the $\mathcal{H}_\rho$ is just $\mathcal{H}$ in disguise:

Proposition 4.8. *Given a net of C^*-algebras satisfying the Haag-Kastler axioms (1)-(6), let ϕ denote a global state (analytic in energy) on the quasilocal algebra $\mathfrak{A}$. If $\rho = \phi|_{\mathfrak{A}(\mathcal{O})}$, then there exists a unitary equivalence between $(\pi_\rho(\mathfrak{A}(\mathcal{O}))'', \mathcal{H}_\rho)$ and $(\pi_\phi(\mathfrak{A}(\mathcal{O}))'', \mathcal{H}_\phi)$.*[31]

It follows from this result that we cannot view the density operator D_ρ as local because it is an element of a certain "local" Hilbert space. To make matters worse, since density operators are compact operators, D_ρ has a spectral decomposition consisting of finite-dimensional projections.[32] Every finite-dimensional projection is a finite projection. Since type III algebras lack finite projections, this means that D_ρ cannot be an element of the local algebra $\pi_\rho(\mathfrak{A}(\mathcal{O}))''$ either. So D_ρ doesn't inherit its locality from being a local observable itself. All that's left for the spacetime state realist to do is bite the bullet and say that the operator $D_\rho \in B(\mathcal{H})$ is local *simply because* it encodes the expectation values for the algebra of observables localized in $\mathcal{O}$ (and remains silent on observables in $\mathcal{O}'$). But this means that the local observables must be conceptually prior to the local state, which is hard to square with the claim that the field of local density operators exhausts the fundamental ontology of the theory.[33]

[31] *Proof.* Let x_ϕ, x_ρ denote the cyclic vectors representing ϕ, ρ in the GNS representations $(\pi_\phi, \mathcal{H}_\phi)$, $(\pi_\rho, \mathcal{H}_\rho)$. First we show that $\mathcal{H}_\rho$ is a closed subspace of $\mathcal{H}$. Taking $\pi_\phi(\mathfrak{A}(\mathcal{O}))$ as a subalgebra of $B(\mathcal{H}_\phi)$, we consider the compression map $\kappa : \pi_\phi(A) \to \pi_\phi(A)|_{\mathcal{K}}$ with $\mathcal{K} = \overline{\pi_\phi(\mathfrak{A}(\mathcal{O}))x_\phi}$, a $\pi_\phi(\mathfrak{A}(\mathcal{O}))$-invariant subspace of $\mathcal{H}_\phi$. Thus we can view $\kappa \circ \pi_\phi$ as a representation of $\mathfrak{A}(\mathcal{O})$ on $\mathcal{K}$. Clearly, x_ϕ is a cyclic vector for $(\kappa \circ \pi_\phi, \mathcal{K})$, moreover $\rho = \phi \circ (\kappa \circ \pi_\phi)$. These two conditions define the local GNS representation $(\pi_\rho, \mathcal{H}_\rho)$ up to unitary equivalence (Kadison and Ringrose, 1997a, Prop. 4.5.3). Therefore $(\pi_\rho, \mathcal{H}_\rho)$ and $(\kappa \circ \pi_\phi, \mathcal{K})$ are unitarily equivalent.

Next, we show that in fact $\mathcal{K} = \mathcal{H}_\phi$. By the Reeh-Schlieder theorem, we obtain $\overline{\pi_\rho(\mathfrak{A}(\mathcal{O}))''x_\phi} = \overline{\pi_\phi(\mathfrak{A}(\mathcal{O}))''x_\phi} = \mathcal{H}$. But if x_ϕ is cyclic for $\pi_\phi(\mathfrak{A}(\mathcal{O}))''$, it must also be cyclic for $\pi_\phi(\mathfrak{A}(\mathcal{O}))$. Therefore $\mathcal{H}_\rho = \overline{\pi_\phi(\mathfrak{A}(\mathcal{O}))x_\phi} = \mathcal{H}_\phi$. □

[32] A linear operator $A \in B(\mathcal{H})$ is *compact* if the set $\{Ax \mid ||x|| \leq 1\}$ is compact in the norm topology. For a full discussion of the spectral theory of compact operators, see Conway (1990, Ch. VII, §7).

[33] Additionally, Baker (2014) argues that interpreting the local state as a summary of expectation values doesn't work well with Wallace's commitment to Everettian QM.

4.3.1 Cutoffs and the Type III Property

Elsewhere (Wallace, 2006, 2011, 2012), Wallace has defended an interpretive view of QFT motivated by the idea of an effective field theory. He argues that we should be wary of ontological conclusions which depend sensitively on the extreme short distance physics because we expect our theories to break down in this regime and be replaced by an eventual theory of quantum gravity. Moreover, the best way we have to understand Yang-Mills type theories like the Standard Model involve imposing a floating UV cutoff. Renormalization group techniques are then used to show that the low energy physics do not depend on the details of this cutoff. This line of thought suggests a different response to the arguments against spacetime state realism presented above — perhaps type III algebras are idealizations which only appear in the absence of a UV cutoff. In a footnote to his discussion of spacetime state realism, Wallace suggests exactly this line of argument:

> In the standard presentations of AQFT, the algebra $\mathfrak{R}(\mathcal{O})$ is infinite-dimensional (even-though it is associated with a spatially finite region); as such it has multiple non-isomorphic representations, and so different states lead to different Hilbert spaces. This makes it less clear that we are licensed to talk about $\mathcal{H}_{\mathcal{O}}$ as 'the' Hilbert space for the region $\mathcal{O}$. On the modern (Wilsonian) understanding of renormalization, however, this is an artefact of the formalism, which disappears when we cut off the unphysically high-energy degrees of freedom; as such we ignore this complication in the text (Wallace, 2012, p. 301, fn. 13)

Here's a physical gloss on Wallace's idea: as we know, fields have an infinite number of degrees of freedom. This comes both from the fact that they are extended across the infinite expanse of Minkowski spacetime and from the fact that they are continuous — field-variables can be assigned to arbitrarily small compact regions. But when we consider the local algebra of observables in a bounded region, $\mathcal{O}$, we can safely ignore the extreme long distance physics, since we're only concerned with local degrees of freedom. Moreover, if a take the effective

field theory picture seriously and impose a floating UV cutoff, then we can safely ignore the extreme short distance physics. As a consequence, inside $\mathcal{O}$ the field can be characterized by finitely many degrees of freedom, and we're back safely in the well-understood territory of finite, type I_n (or possibly II_1) von Neumann algebras.

The argument is simple, physically motivated, and correct if applied to classical fields. In the quantum domain, however, entanglement makes it impossible to ignore long-distance effects, even if a short-distance cutoff is imposed. As we saw in Chapter 1, the Reeh-Schlieder theorem survives if we assume a cutoff version of the weak additivity axiom. One can use this fact to prove that the local algebras $\mathfrak{R}(\mathcal{O})$ must be infinite.

Proposition 4.9. *Assume Haag-Kastler axioms (1)-(6*). It follows that for any bounded region $\mathcal{O} \gg \mathcal{O}_\lambda$, with UV volume cutoff λ, the local von Neumann algebras $\mathfrak{R}(\mathcal{O})$ in any sector analytic for the energy must be infinite (i.e. type I_n, II_∞, or III).*[34]

Thus we find that in the vacuum sector (or any sector analytic for the energy) the local degrees of freedom are so strongly entangled with the global degrees of freedom, that we cannot idealize them away. The local algebras are always of infinite type. This conclusion depends on the long distance physics, not the short distance physics, and hence survives even if we view QFT as an effective field theory.

This is not yet enough to secure the type III property. Fredenhagen's proof that the double-cone algebras, $\mathfrak{R}(\mathcal{D})$, must be isomorphic to the unique hyperfinite type III_1 factor relies on the fact that the algebras become conformally invariant in the UV scaling limit. It is therefore not obvious that the conclusion would continue to hold in the case of effective field theories. While we do not have an elementary proof of the type III property from cutoff-safe

[34] *Proof.* From the spectrum condition and cutoff weak additivity we obtain the Reeh-Schlieder theorem for local algebras above the cutoff (see Appendix C). Thus the vacuum state is cyclic and separating for both $\mathfrak{R}(\mathcal{O})$ and $\mathfrak{R}(\mathcal{O}')$ and we can apply the results of Tomita-Takesaki modular theory. Since $\mathcal{O}'$ is unbounded, it contains a wedge region. Wedge algebras are type III_1 by Thm 4.2, so isotony entails that $\mathfrak{R}(\mathcal{O}')$ must be infinite. By microcausality $\mathfrak{R}(\mathcal{O}') \subseteq \mathfrak{R}(\mathcal{O})'$, so $\mathfrak{R}(\mathcal{O})'$ is infinite. But the Tomita-Takesaki theorem establishes that $J\mathfrak{R}(\mathcal{O})J = \mathfrak{R}(\mathcal{O})'$ is an anti-automorphism, so $\mathfrak{R}(\mathcal{O})$ is infinite. □

assumptions at this stage, the following argument strongly supports the conclusion that at the very least, the local algebras cannot be type I_∞.

First, we must introduce a new independence condition called *strict locality*.[35] Let $\mathfrak{M}$, $\mathfrak{N}$ be two commuting von Neumann algebras on a Hilbert space $\mathcal{H}$. The ordered-pair $(\mathfrak{M}, \mathfrak{N})$ is strictly local if for any projection operator $E \in \mathfrak{M}$ and any unit vector $x \in \mathcal{H}$, there is a unit vector $y \in E\mathcal{H}$ such that $\langle x, Ax\rangle = \langle y, Ay\rangle$ for all $A \in \mathfrak{N}$. From a physical standpoint, let $\mathfrak{M}$, $\mathfrak{N}$ be local algebras for spacelike separated regions $\mathcal{O}_1$, $\mathcal{O}_2$. The vectors in $E\mathcal{H}$ represent states such that $\phi(E) = 1$, i.e. states which have the property represented by the projection operator E. Strict locality requires that for any global state, we cannot glean information about the outcome of a measurement of E in $\mathcal{O}_1$ by examining results of measurements localized in $\mathcal{O}_2$.

If the local algebras are type I_∞ and duality holds, then each local algebra satisfies strict locality with respect to the algebra associated with the life-preserver region around it:

Proposition 4.10. *If the local algebra $\mathfrak{R}(\mathcal{D})$ is type I and duality holds for double-cones in the vacuum sector, then the pair $(\mathfrak{R}(\mathcal{D}), \mathfrak{R}(\mathcal{L}))$ satisfy strict locality.*[36]

But if we inflate the life-preserver, and take the limit $\mathcal{L} \to \mathcal{D}'$, we hit a snag. Assuming the local algebras are factors, then $\mathfrak{R}(\mathcal{D})$ satisfies strict locality with respect to $\mathfrak{R}(\mathcal{D}') = \mathfrak{R}(\mathcal{D})'$ iff it is type III.[37] So there must be a discontinuity in the limit $\mathcal{L} \to \mathcal{D}'$.

While mathematically consistent, this state of affairs is physically rather puzzling. Under the assumption that the local algebras are type I factors, the local observables are strictly local with respect to observables in the nearby causal complement, $\mathcal{L}$. This is a reasonably

[35] There are actually two different, mathematically independent conditions that go by this name in the literature. Here we mean the stronger one, what Summers (1990) calls *strict locality in the vector sense*. The two conditions are equivalent, however, if the algebras in question have a separating vector. This fact proves key in the subsequent argument.

[36] *Proof.* Since $\mathfrak{R}(\mathcal{D})$ is type I, it trivially satisfies the split property. In particular, there is an isomorphism $W : \mathcal{H} \to \mathcal{H}\otimes\mathcal{H}$ such that $W^*\mathfrak{R}(\mathcal{D})W = B(\mathcal{H})\otimes I$ and $W^*\mathfrak{R}(\mathcal{D})'W = I\otimes B(\mathcal{H})$. As $\mathfrak{R}(\mathcal{L}) \subset \mathfrak{R}(\mathcal{D}') = \mathfrak{R}(\mathcal{D})'$, $\mathfrak{R}(\mathcal{D})$ and $\mathfrak{R}(\mathcal{L})$ are also split. The joint algebra $\mathfrak{R}(\mathcal{D}) \vee \mathfrak{R}(\mathcal{L})$ is a local algebra (a subalgebra of $\mathfrak{R}(\tilde{\mathcal{D}})$), therefore by the Reeh-Schlieder theorem, the restriction of Ω is cyclic and separating. But since the joint algebra has a separating vector, the split property entails strict locality (Summers, 1990, Lem. 4.4). □

[37] See Summers (1990, Cor. 4.6).

strong independence condition, indicating that we can disentangle the local system from nearby vacuum fluctuations. But if we can disentangle a system from nearby vacuum fluctuations, we should be able to disentangle it from distant fluctuations. So it comes as a shock that strict locality fails in the limit as $\mathcal{L}$ expands to cover $\mathcal{D}'$. If the local algebras are type III, this pathology vanishes. It is incumbent on the defender of type I algebras then, to provide a physical explanation for why strict locality fails in the long distance limit.

One might wonder about the physical motivation for duality. Microcausality requires that $\mathfrak{R}(\mathcal{D}') \subset \mathfrak{R}(\mathcal{D})'$. If it is possible to add additional elements to $\mathfrak{R}(\mathcal{D})$ while preserving microcausality, then $\mathfrak{R}(\mathcal{D}')$ grows larger while $\mathfrak{R}(\mathcal{D})'$ shrinks. Therefore duality amounts to the assumption that the double-cone algebras are maximal. There are two primary lines of support for this maximality assumption. The first comes from transitivity of locality arguments found in Borchers and Yngvason (1994). The second comes from the Bisognano-Wichmann theorem. Despite this, duality is expected to fail in certain physically motivated cases. In the vacuum sector, duality is equivalent to the absence of spontaneous symmetry breaking. In charged sectors, duality holds iff the gauge group is abelian. In both cases we expect weaker forms of duality to hold (e.g. *essential duality* or *wedge duality*), but it is not obvious how to employ these conditions in the argument above.

For the defender of spacetime state realism, however, rejecting duality looks like a dead end. Under additional natural assumptions, if the local algebras are type I factors, then duality must hold.[38]Furthermore, Wallace and Timpson clearly intend their view to be a highly general ontological picture for quantum theory. Since the type III property holds generically in cutoff-free models, a substantial amount of ground already has to be ceded. It would be problematic if the viability of spacetime state realism hinges on contingent model-theoretic details like the failure of duality.

[38]These assumptions are *strong additivity*, *outer continuity* and *primitive causality*. See Kraus (1964) as well as Horuzhy (1990, Prop. 1.3.6). Assuming outer continuity we can write $\mathcal{D} \cup \mathcal{D}'$ as the projective limit of spacelike slices. Given primitive causality and strong additivity we then have $\mathfrak{R}(\mathcal{D} \cup \mathcal{D}') = \mathfrak{R}(\mathcal{D}) \vee \mathfrak{R}(\mathcal{D}') = B(\mathcal{H})$. If the local algebras are type I factors, then they are *normal*, i.e. every von Neumann subalgebra is equal to its own double commutant. Therefore by microcausality, we have $\mathfrak{R}(\mathcal{D}') = (\mathfrak{R}(\mathcal{D}')' \cap \mathfrak{R}(\mathcal{D})')' \cap \mathfrak{R}(\mathcal{D})' = B(\mathcal{H}) \cap \mathfrak{R}(\mathcal{D})' = \mathfrak{R}(\mathcal{D})'$.

It should also be stressed that many of the arguments in the previous section do not depend on the full strength of the type III property. For example, proposition 4.8 only requires the Reeh-Schlieder theorem. Thus assuming the standard Haag-Kastler axioms plus cutoff weak additivity, there is no hope of resurrecting the local GNS Hilbert space via appeal to type I_∞ or II_∞ algebras. Additionally, it turns out that we can prove that $\mathfrak{R}(O)$ does not contain (true) density operators from cutoff-safe assumptions, without requiring the type III property.

Proposition 4.11. *Assume Haag-Kastler axioms (1)-(6*). Then $\mathfrak{R}(O)$ cannot contain any compact operators.*[39]

Therefore the local density operators cannot be local observables regardless of how the type debate plays out.[40]

4.3.2 Split Localization to the Rescue?

If left unchecked, the type III property causes serious problems for spacetime state realism. Moreover the results of the previous section suggest that these consequences cannot be easily evaded by appealing to effective field theories. Can spacetime state realism deploy the split condition to salvage the view? Even though proposition 4.7 ensures that we still won't be able to disentangle $\mathfrak{R}(\mathcal{D})$ and $\mathfrak{R}(\mathcal{D}')$, if the split property holds, then we can always find a slightly larger region $\tilde{\mathcal{D}}$, such that $\mathfrak{R}(\mathcal{D})$ and $\mathfrak{R}(\tilde{\mathcal{D}}')$ can be disentangled. Thus there exists

[39] *Proof.* Theorem 4.2 establishes that $\mathfrak{R}(\mathcal{W})$ is a type III_1 factor. (The theorem requires Haag-Kastler axioms (1)-(6), but the only use of weak additivity is to prove that the Reeh-Schleider theorem holds, therefore we can safely replace the assumptions with (1)-(6*).) Any bounded $\mathcal{O}$ can be localized in the interior of a wedge, $\mathcal{W}$. If $\mathfrak{R}(\mathcal{O})$ contains compact operators, then it contains finite-dimensional projections, namely the spectral decomposition of the compact operators. Hence $\mathfrak{R}(\mathcal{W})$ does as well. Let P be one such projection. The finite-dimensional subspace $P\mathcal{H}$ is invariant under $\mathfrak{R}(\mathcal{W})'$, which allows us to define the *induced von Neumann algebra*, $\mathfrak{R}(\mathcal{W})'_P$, by restricting $A \in \mathfrak{R}(\mathcal{W})'$ to $P\mathcal{H}$. The induced mapping, $\iota : A \to A|_{P\mathcal{H}}$, defines an isomorphism between $\mathfrak{R}(\mathcal{W})'$ and $\mathfrak{R}(\mathcal{W})'_P$ (since $\iota(P) = I$), but the former is type III_1 while the latter is necessarily finite. □

[40] There is a loophole here, if the local algebras are type I_∞ or type II_∞, then even though they will not contain true density operators, there could still be local operators that "look like" density operators from the perspective of the local algebra — an operator D such that $\phi(A) = \mathrm{Tr}(DA)$, for every ϕ. Here Tr is the semi-finite, rather than finite trace. Type III algebras lack a trace of any kind.

an isomorphism $W : \mathcal{H} \to \mathcal{K}_1 \otimes \mathcal{K}_2$ such that $\mathfrak{R}(\mathcal{D})$ is isomorphic to a subalgebra of $B(\mathcal{K}_1) \otimes I$ and $\mathfrak{R}(\tilde{\mathcal{D}}')$ is isomorphic to a subalgebra of $I \otimes B(\mathcal{K}_2)$. Is it possible to interpret $\mathcal{K}_1$ as the long-sought local Hilbert space?

Not really. The trouble is that we cannot define the split uniquely without the canonical interpolating type I factor, and $\mathfrak{N}_\Lambda$ is not a local algebra. There is no region of spacetime that we can assign $\mathcal{K}_1$ to such that $B(\mathcal{K}_1)$ represent the observables localized therein. The split condition *can* help us isolate subsystems in QFT, but it requires detailed information about the local algebras of observables to be used effectively. Given a nested set of local algebras, their joint modular structure generates $\mathfrak{N}_\Lambda$. But this is not information that the spacetime state realist can avail himself of directly. The essentially finite dimensional subspace, $\mathcal{N}_{\beta,r} \in \mathcal{H}$ (see Eq. 4.5), defined with the help of the nuclearity condition, is of no help either. It's too small to be a Hilbert space itself, and once we take the closed linear span, we find that $\overline{\mathcal{N}_{\beta,r}} = \mathcal{H}$.[41]

Thus we find that because of the type III property, describing the localization of states in QFT requires ineliminable reference to localized observables. If we want to give an ontology for QFT that is even approximately local, as the spacetime state realist does, we will need to rely on algebraic tools. But this requires that we make room in our fundamental ontology for localized quantities in addition to states.

This conclusion is in broad agreement with other criticisms of Wallace and Timpson's view. For example, Baker (2014) points out that since the local algebras are infinite, their associated GNS Hilbert spaces will be separable, and thus isomorphic as Hilbert spaces. The corresponding collections of local density operators (which are just probability measures over

[41] As noted briefly in §4.2 (footnote 17), the collection of canonical type I factors has the structure of a Poincaré covariant net. The statespace realist might attempt to exploit this structure to define the network of type I factors directly, without appeal to the type III local algebras. The basic idea would be to assign a Hilbert space to each pair of nested double-cones such that bounded operators on that space correspond to the canonical type I factor that would be obtained had we started with the corresponding type III local algebras. Since the type I factors are necessarily infinite by Prop. 4.9, these Hilbert spaces will all be separable, hence isomorphic. A privileged set of embedding maps will be required to recover the structure of the semilocal type I net. It is far from obvious that such a construction is possible, much less if it can recover all of the information in the original net of local algebras. There are however a number of open technical questions in this direction which are of independent interest.

1-dimensional subspaces of the local Hilbert spaces) will be similarly isomorphic. Thus if we ignore the local algebras, all there is to a QFT is an assignment of isomorphic Hilbert spaces and local density operators to regions. This leads to absurdity. All QFTs on the same background spacetime (including free and interacting theories) will turn out to be physically equivalent. Baker concludes that something must be wrong with Wallace's argument for the equivalence of algebraic localization and Hilbert space quotienting. Our discussion here helps to explain why. Since local algebras in QFT are type III, we cannot use a naive spatiotemporal partition to quotient the global Hilbert space. The only quotienting procedures available factor through the local algebras themselves.

Chapter 5

Can Quantum Thermodynamics Save Time?

In both classical and quantum theories defined on a fixed background spacetime manifold, the physical flow of time is represented in much the same way. Time translations correspond to a continuous 1-parameter subgroup of spacetime symmetries, and the dynamics are implemented either as a parametrized flow on statespace (Schödinger picture) or a parametrized group of automorphisms of the algebra of observables (Heisenberg picture). In generally covariant theories, such as General Relativity or background independent theories of quantum gravity, this picture breaks down — there is no canonical time-translation subgroup at the global level, nor is there a clear way to represent dynamics locally in terms of the Schrödinger or Heisenberg pictures.[1] Moreover, without a preferred flow on the space of states representing time, the standard way to represent physical change via functions on this space that take on different values at different times, also fails. Where and how to locate time and change in such theories poses a deep interpretational challenge for an adequate account of gravitational physics. The *problem of time*, as it is widely known, has been discussed extensively both

[1]Here, by general covariance we mean *strong* general covariance. This requires that not only is Diff($\mathcal{M}$) a group of symmetries for the equations of motion of the theory, but that it is a set of gauge symmetries, meaning that diffeomorphism-liked solutions are identified as physically equivalent. See Belot (2005) §7 for a further discussion.

by physicists researching the foundations of quantum gravity, as well as by philosophers of physics working on relativity.

While the problem has many faces (it's probably more accurate to refer to the *problems* of time), two will be of central concern here. The first is the issue of selecting a preferred time flow from the multitude of possible candidates that present themselves even in the most well-behaved spacetimes. The second is the explanatory challenge of understanding this flow in purely gauge invariant terms. Ultimately, both problems stem from treating the diffeomorphism group of the spacetime manifold as a group of gauge symmetries, identifying any two diffeomorphism-linked solutions to the equations of motion as representations of the same possible world. In the case of GR, we can restrict attention to globally hyperbolic solutions of the Einstein Field Equations for which there exist global time functions whose $t =$const surfaces are cauchy. Even for these temporally well-behaved spacetimes, there will typically be numerous potential time functions, and due to general covariance, no way to select a preferred physical time from among them. Moreover, none of these functions will have the feature that gauge invariant physical quantities take on different values at different times. If the standard connection between gauge symmetries and ontology is assumed, then it is not clear how any of these candidate flows can describe genuine physical change. The dynamics of the theory appear to be frozen.

Connes and Rovelli (1994) proposes a radical solution to the problem: the flow of time (not just its direction) has a thermodynamic origin. In a non-generally covariant theory, an equilibrium state is defined with respect to a background time flow (e.g. dynamical stability and passivity constraints require a group of time translations). Conversely, given an equilibrium state one can derive the time flow from the definition of the Gibbs canonical ensemble (or more generally, the KMS condition). Rovelli (2011) exploits this converse connection, arguing that in a generally covariant theory, *any* statistical state defines a notion of time according to which it is an equilibrium state, and that this state-dependent *thermal time* should be identified with physical time.

Drawing upon tools from modular theory, Connes and Rovelli demonstrate how this idea can be rigorously implemented in the case of generally covariant quantum theories. Minimally, such a theory can be thought as a non-commutative C^*-algebra of diffeomorphism-invariant observables, $\mathfrak{A}$, along with a set of physically possible (normal) states.[2] As we have seen in chapter 2, given any such state, the Tomita-Takesaki theorem defines a canonical 1-parameter group of automorphisms — the modular automorphism group — on the GNS von Neumann algebra, $\pi(\mathfrak{A})''$. Moreover, the state satisfies the KMS condition with respect to the modular dynamics for inverse temperature $\beta = 1$. The *thermal time hypothesis* (TTH) identifies physical time with this state-dependent modular flow.

The idea is an intriguing one that, to date, has received little attention from philosophers.[3] This chapter represents a modest initial attempt to sally forth into uncharted yet potentially rich philosophical territory. In the next section we review the details of the TTH and its quantum implementation via modular theory.[4]

5.1 The Thermal Time Hypothesis

Lying behind Connes and Rovelli's proposal, is a very general framework for constructing a timeless theory of mechanics. Ordinarily, mechanics is viewed as consisting of a kinematical specification of a set of possible states and physical quantities (observables), and a dynamical specification of how these states and quantities evolve through time. Rovelli (2011) advocates replacing this picture with a more general one that conceives of mechanics as a theory of relative correlations between physical quantities. Rovelli's picture requires an important distinction between *partial* and *full* observables. Partial observables are quantities that physical measuring devices can be responsive to, but whose (expectation) value cannot be predicted given the state alone. A full observable is understood as a coincidence or correlation

[2]See Brunetti et al. (2003) for a formal development of this basic idea using category theory.

[3]Earman (2002) and Earman (2011) are notable exceptions.

[4]Material from this chapter was originally presented at the *IPP Conference on the Mathematical and Conceptual Foundations of Physics*, University of Pittsburgh, April 2013.

of partial observables whose (expectation) value can be predicted given the state. In GR, an example of a partial observable is the proper time measured by a clock carried along an observer's worldline.[5] A full observable might be given by the value of the observer's proper time at a point where his world line intersects with the worldline of some other observer. (In this sense, full observables can be thought of as a generalization of Einstein's coincidence observables.) Only measurements of full observables can be directly compared to the predictions made by the mechanical theory.

A timeless mechanical system is given by a triple $(\mathcal{C}, \Gamma, f)$. $\mathcal{C}$ is the configuration space of partial observables, q^a. A *motion* of the system is given by an unparametrized curve in $\mathcal{C}$, representing a sequence of correlations between partial observables. The space of motions, Γ is the statespace of the system and is typically presymplectic. The evolution equation is given by $f = 0$, where f is a map $f : \Gamma \times \mathcal{C} \to V$, and V is a vector space. For systems that can be modeled using Hamiltonian mechanics, Γ and f are completely determined by a surface Σ in the space $T^*\mathcal{C}$ of partial observables and their conjugate momenta p_a.[6] This surface can be defined by the vanishing of the Hamiltonian function $H : T^*\mathcal{C} \to \mathbb{R}$.

If the system has a preferred external time variable, the Hamiltonian can be decomposed as

$$H = p_t + H_0(q^i, p_i, t) \tag{5.1}$$

where t is one of the partial observables in $\mathcal{C}$ that corresponds to time. (The quantity $E = -p_t$ is the energy.) In a generally covariant mechanical system, however, H treats all partial observables on the same footing, hence there is no canonical decomposition of the form (1). At a fundamental mechanical level, a generally covariant theory is timeless, although it remains possible for a notion of time to emerge in some limiting regime. The

[5]The proper time τ is defined as the length of the observers worldline as determined by the metric g_{ab}. Given an interval $I \in \mathbb{R}$, if $\gamma : I \to \mathcal{M}$ is a timelike curve with tangent vector field ξ^a, the line element is given by

$$\tau = \int_I \sqrt{g_{ab}\xi^a\xi^b}ds.$$

[6]$T^*\mathcal{C}$ is the cotangent bundle over the configuration space of partial observables.

basic idea behind the TTH is that a coarse grained statistical description of the system might single out a preferred time variable. The motivation for this thought is the following observation: a closed system left to thermalize will eventually settle into a time-independent equilibrium state. Viewed as part of a definition of equilibrium, this thermalization principle requires an antecedent notion of time. Rovelli's idea is to invert this definition and use the notion of an equilibrium state to select a partial observable in $\mathcal{C}$ as time.

Three hurdles present themselves. The first is a coherent mathematical characterization of equilibrium states. The second challenge is to find a method for extracting information about the associated time flow from a specification of the state. Finally, in order to count as an emergent explanation of time, one has to show that the partial observable selected behaves as a traditional time variable in the appropriate non-generally covariant limit. In the context of generally covariant QFT, the Tomita-Takesaki theorem gives Connes and Rovelli a way to surmount the first two hurdles. The third hurdle proves more difficult, although the Bisognano-Wichmann theorem presents a possible way forward.

As discussed in §2.1.1, the KMS condition provides a generalization of the idea of an equilibrium state for quantum systems with infinite degrees of freedom. Recall that KMS states are stable, passive, and invariant under the dynamics. Moreover in the finite limit, the KMS condition reduces to the standard Gibbs postulate. The problem for the TTH is that the KMS condition is defined relative to a chosen background dynamics. A state, ρ, on a von Neumann algebra, $\mathfrak{M}$, satisfies the KMS condition for inverse temperature $0 < \beta < \infty$ with respect to a 1-parameter group of automorphisms, $\{\alpha_t\}$, of $\mathfrak{M}$ if for any $A, B \in R$ there exists a complex function $F_{A,B}(z)$, analytic on the strip $\{z \in \mathbb{C} | 0 < \mathrm{Im} z < \beta\}$ and continuous on the boundary of the strip, such that

$$\begin{aligned} F_{A,B}(t) &= \rho(\alpha_t(A)B) \\ F_{A,B}(t + i\beta) &= \rho(B\alpha_t(A)) \end{aligned} \tag{5.2}$$

for all $t \in \mathbb{R}$. Since the ultimate goal is to identify a privileged group of automorphisms implementing the dynamics, we cannot begin with KMS states unless we have an alternate means of characterizing them.

Connes and Rovelli propose a complete end-run around this problem. Modular theory tells us that *any* faithful normal state defines a privileged group of automorphisms according to which it is a KMS state. Let ϕ be a faithful, normal state over the C^*-algebra $\mathfrak{A}$, representing the observables of a generally covariant quantum theory. In the corresponding GNS representation, ϕ is represented by a cyclic, separating vector $\Phi \in \mathcal{H}_\phi$. By the Tomita-Takesaki theorem, the associated unitary operators $\{\Delta^{it} | t \in \mathbb{R}\}$ form a strongly continuous unitary group, inducing a 1-parameter automorphism group on the GNS von Neumann algebra, $\sigma_t(A) := \Delta^{it} A \Delta^{-it}$, for all $A \in \pi(\mathfrak{A})''$, $t \in \mathbb{R}$. Since $\Delta\Phi = \Phi$, the state is invariant under the flow of the modular automorphism group, $\phi(\sigma_t(A)) = \phi(A)$. Furthermore, $\phi(\sigma_t(A)B) = \phi(B\sigma_{t-i}(A))$. Thus ϕ satisfies the KMS condition relative to $\{\sigma_t\}$ for inverse temperature $\beta = 1$.[7]

For any faithful, normal state, this procedure identifies a partial observable, the thermal time, $t_\phi := t$, parametrizing the flow of the (unbounded) thermal hamiltonian $H_\phi := -\ln \Delta$, which has Φ as an eigenvector with eigenvalue zero. We can then go on to decompose the timeless Hamiltonian $H = p_{t_\phi} + H_\phi$. Associated with any such state, there is a natural "flow of time" according to which the system is in equilibrium. But in what sense does this thermal time flow correspond to various notions of physical time? In particular, how is thermal time related to the proper time measured by a localized observer? At present, a general theorem linking thermal time and proper time is out of the question since we lack an

[7]The mysterious temperature $\beta = 1$ is in fact an arbitrary convention that can be removed by rescaling the temperature variable. A state is a $\beta = 1$ KMS state with respect to modular automorphisms σ^t iff it is an arbitrary β KMS state with respect to σ^s, where $s = t/\beta$.

axiomatic formulation of AQFT in curved spacetime.[8] In the interim, a number of intriguing connections have been obtained in special cases.

For a constantly accelerating, immortal observer in Minkowski spacetime, the region causally connected to the points comprising the observer's worldline is the Rindler wedge. In standard coordinates we can explicitly write the observer's trajectory as

$$\begin{aligned} x^0(\tau) &= a^{-1}\sinh(\tau) \\ x^1(\tau) &= a^{-1}\cosh(\tau) \\ x^2(\tau) &= x^3(\tau) = 0, \end{aligned} \tag{5.3}$$

where τ is the observer's proper time. The wedge region is defined by the condition $x^1 > |x^0|$. This represents the region of spacetime that the observer has informational access to via localized experiments. The Bisognano-Wichmann theorem then tells us that in the vacuum sector, the modular automorphism group for the wedge implements wedge preserving Lorentz boosts — Δ^{it} is given by the boost $U(t) = e^{2\pi i t K_1}$ (where K_1 is the representation of the generator of an x^1-boost). Additionally, since the Lorentz boost $\lambda_W(a\tau)$ implements a proper time translation along the orbit of an observer with acceleration a, the operator $U(\tau) = e^{ai\tau K_1}$ can be viewed as generating evolution in proper time. Comparing these two operators, we find that the observer's proper time is directly proportional to the thermal time,

$$t = \frac{2\pi}{a}\tau. \tag{5.4}$$

Note that the Unruh temperature measured by the observer is $T = a/2\pi k_b = 1/k_b\beta$, this leads Connes and Rovelli to propose that the temperature can be interpreted as the ratio between thermal and proper time. Moreover, not only does the thermal time coincide with

[8]Since generic solutions of the Einstein equations lack isometries, a number of the axioms of AQFT need to be reformulated including covariance and the spectrum condition. Modular theory may be of great service here as well. An important test for the TTH is to extend the analysis here to generally covariant contexts like those discussed by Brunetti et al. (2003) and Fredenhagen (2004).

the (rescaled) proper time along the orbits of constant acceleration, but if an observer extends her proper time to a global time for the wedge via the process of Einstein synchronization, this global time continues to coincide with the thermal time flow. Hence, the dynamics given by the flow of thermal time align with the dynamics a constantly accelerating observer would naturally employ to describe the entire wedge region.[9]

Similar results have been obtained for the local algebras associated with double-cones and forward lightcones in conformally invariant QFTs with spacetime $d > 2$.[10] In these cases the modular operators associated with the local algebras in the vacuum sector have a geometric interpretation similar to the Rindler wedge (since all three regions can be transformed into one another by order-preserving conformal transformations). Double-cones are of particular interest, since these correspond to regions causally connected to mortal observers like us. Moreover, the results of any local experiment can be modeled via the restriction of the global algebra to a bounded region contained inside a double-cone. The relationship between thermal time and proper time in this case is more complicated, due to the fact that the the proper time experienced by a finite observer is bounded while the modular time is unbounded. As before we have a functional relation, $t = t(\tau)$, but now

$$t(\tau) = \frac{2\pi}{La^2}(\sqrt{1 + a^2L^2} - \cosh a\tau) \ , \tag{5.5}$$

[9]For a given uniformly accelerating observer with acceleration a, we can rewrite the Minkowksi metric in *Rindler coordinates*, ξ, η,

$$x^0 = \frac{1}{a}e^{a\xi}\sinh(a\eta), \quad x^1 = \frac{1}{a}e^{a\xi}\cosh(a\eta),$$

for $x^1 > |x^0|$. The metric then takes the form,

$$ds^2 = e^{2a\xi}(-d\eta^2 + d\xi^2).$$

The observer moves on trajectories $\eta = \tau$ and $\xi = 0$, so τ extends to a coordinate time across the wedge with surfaces $\eta =$const as simultaneity hypersurfaces. Since the metric is independent of η, the infinitesimal proper time translations ∂_η are a Killing field. (In standard coordinates this is equivalent to the Killing field $a(x^1\partial_{x^0} + x^0\partial_{x^1})$ corresponding to a boost in the x^1-direction). The thermal time coordinates are just a constant rescaling of this coordinate system by $a/2\pi$.

[10]Martinetti and Rovelli (2003)

where L is the lifetime of the observer. For most of the observer's lifespan, $\beta(\tau)$ is approximately constant, allowing the Unruh temperature to be interpreted as the local ration between thermal and proper time for such observers.[11] As in the Rindler case, the thermal time agrees with the global time for the diamond region determined by the Einstein synchronization procedure conducted by uniformly accelerating observers.

We can summarize the main content of the TTH with the following:

Thermal Time Hypothesis (Rovelli-Connes)**.** *In a generally covariant quantum theory, the flow of time is defined by the state-dependent modular automorphism group. The Unruh temperature measured by an accelerating observer represents the ratio between this time and her proper time.*

The TTH has three broad components: (a) the general idea that the flow of time is selected at the level of statistical mechanics in a fundamentally timeless, generally covariant theory, (b) a quantum model for such a selection mechanism identifying thermal time with the state-dependent modular flow on the algebra of observables, and (c) a conjecture that in the limit where a geometric notion of proper time exists, the Unruh temperature is interpretable as the ratio of thermal time to proper time. This is a bold idea with a numerous potential implications for quantum physics and cosmology which require further study. Over the next three sections, we will consider a series of technical and philosophical challenges for the TTH.

5.2 Thermal Time and Proper Time

Much of the theoretical support for the TTH comes from the close connection between thermal time and proper time established by the Bisognano Wichmann theorem. But unless our world has fundamental conformal symmetry, the theorem only applies to wedge algebras

[11] Note that unlike in the Rindler case, the temperature seen by non-accelerating observer does not vanish. If $a = 0$, there is a residual temperature $T_D = \frac{2}{\pi L}$. This raises question of whether or not a finite observer might determine the date of her death by careful measurements of T_D. Martinetti (2004) has shown that the Heisenberg uncertainty principle prevents this. A finite observer will not live long enough to determine T_D with the required accuracy.

in the vacuum sector. This limited domain of applicability makes it hard to see how to extend the connection to a broader class of more physically realistic observers and states. [12]

We can attempt to generalize these results in the absence of additional conformal symmetry by considering (1) non-eternal observers, (2) non-uniformly accelerating observers, and (3) non-vacuum states. In the case of an eternal, uniformly accelerating observer there are two geometric time flows, thermal and proper, that agree up to a scale factor. The problem for the TTH is that in cases (1)-(3), if the modular automorphism group acts geometrically at all, there are two competing flows which are not related by a simple rescaling. The defender of the TTH must explain why it is thermal time rather than proper time which represents the appropriate physical time for the local observer.

For observers with finite lifespans (regardless of their acceleration), the region of spacetime causally accessible to them is a double-cone formed by the intersection of their future lightcone at birth and their past lightcone at death. Unfortunately, at present there are no general results ensuring that the modular operators of doublecones have a geometric interpretation in the vacuum sector. Yngvason (1994) provides several counterexamples, although they do not satisfy Lorenz covariance. The best systematic results are due to Trebels (1997).[13] We say that a unitarily implemented automorphism $U\mathfrak{R}(\mathcal{O})U^* = \mathfrak{R}(\mathcal{O})$ is *geometric, causal, and order preserving*, if there exists a 1-1 map, $g : \mathcal{O} \to \mathcal{O}$, such that

(i) if $x \in \mathcal{O}$, then $x_g \in \mathcal{O}$ and $x_{g^{-1}} \in \mathcal{O}$,

(ii) if $x, y \in \mathcal{O}$ and $x - y$ is spacelike, then $x_g - y_g$ and $x_{g^{-1}} - y_{g^{-1}}$ are spacelike,

(iii) if $x - y$ is in the forward lightcone, then $x_g - y_g$ and $x_{g^{-1}} - y_{g^{-1}}$ are in the forward lightcone,

(iv) for every $\mathcal{O}^1 \subset \mathcal{O}$, $U\mathfrak{R}(\mathcal{O}^1)U^* = \mathfrak{R}(\mathcal{O}^1_g)$.

[12] Paetz (2010) presses this point, arguing that the examples from the previous section are too special to draw universal conclusions regarding the origins of time.

[13] For a detailed summary of Trebels's thesis work, see Borchers (2000, §3.4).

If the modular automorphism group has a local dynamical interpretation, σ^t is a geometric, causal, order preserving automorphism, for every $t \in \mathbb{R}$. This will give rise to a corresponding group of mappings $g(t)$ satisfying conditions (i)-(iv) above. Using this definition, Trebels goes on to prove the following result for doublecone algebras:

Theorem 5.1 (Trebels)**.** *For every strongly continuous group of geometric, causal, order preserving automorphisms of $\mathfrak{R}(\mathcal{D})$ in the vacuum sector, $g(t)$ is of the form $\exp(mGt)$, where G is the generator of the Hislop-Longo transformations.*

The Hislop-Longo transformations express the geometric action of the doublecone modular operators in a conformal field theory.[14] Martinetti and Rovelli used this expression to derive equation (5.5), expressing the thermal time flow in doublecone regions in such a theory. Hence Trebels's theorem shows that if the modular operators in a non-conformal theory act geometric, causal, and order preserving, they can only act as scaled versions of the modular operators in a conformal theory (i.e. as Poincaré transformations extended by dilations). Thus the conformal case discussed at the end of §5.1 is the best we can hope for.

Given this fact, one can argue quite generally that thermal time and proper time cannot agree exactly for an observer confined to the doublecone. The reason is that the global timelike Killing field $(\partial_s)^\mu$ on Minkowski spacetime does not restrict to a local Killing field on the double-cone. (Nor is it a conformal Killing field.) In contrast, if the modular automorphism group, σ^t acts geometric, causal, and order preserving on the doublecone, then it generates a local (conformal) Killing field, $(\partial_t)^\mu$. One might be able to press further and argue on these general grounds that the thermal Hamiltonian therefore represents a more natural representation of the local dynamics for an observer in the double-cone than the

[14]Hislop and Longo (1982) first computed the modular automorphism group of a conformal diamond. Letting $\mathcal{D} = \{x \mid x < |x^0| + ||\vec{x}||\}$, they define $x^\pm := |x^0| \pm ||\vec{x}||$. The modular group of $\mathfrak{R}(\mathcal{D})$ in the vacuum sector is then given by the geometric formula:

$$x^\pm(\lambda) = \frac{-(1-x^\pm) + e^{-2\pi\lambda}(1+x^\pm)}{(1-x^\pm) + e^{-2\pi\lambda}(1+x^\pm)}.$$

Such transformations are known as Hislop-Longo transformations.

global translations. Moreover, for most of the observer's life, the flow of thermal time and proper time agree to close approximation. Of course, all of this hinges on the local modular generators acting geometrically as they do in the conformal case. Saffary (2005) argues that in massive theories, we should not expect the double-cone modular operators to carry geometric significance, but the formal results backing this conjecture are only partial at this stage.[15]

The case of non-uniformly accelerating observers generates a second set of conceptual worries. Assuming that the observer is eternal, she is confined to the Rindler wedge. In the vacuum sector, her sense of thermal time will be given by the constant flow of the wedge modular automorphism group expressed by the Bisognano-Wichmann relations. Her experience of proper time will fluctuate with her acceleration, however, and it is not clear why it would be natural for her to describe the physical evolution of observables in the wedge using thermal time. Work on the Unruh effect for non-uniformly accelerating observers, indicates that such observers experience an acceleration-dependent thermal bath.[16] This stands in line with predictions of the TTH — the acceleration-dependent temperature reflects the shifting ratio between constant thermal time and acceleration-dependent proper time. The problem is that even if there's a sense in which the thermal time coordinates are more natural on the wedge region, the TTH still has to explain the phenomenological experience of the observer who will presumably age according to her proper time, not the background thermal time flow.

[15] In the massless case, the modular generators are ordinary differential operators, δ_0, of order 1. In the massive case, it has been conjectured that the modular generators are pseudo-differential operators $\delta_m = \delta_0 + \delta_r$, where the leading term is given by the massless generator δ_0 and δ_r is a pseudo-differential operators of order < 1. This second term is thought to give rise to non-local action without geometric interpretation.

[16] Using techniques originally developed for dealing with the Hawking radiation generated by black holes with a time-dependent mass, Jian-yang et al. (1995) have shown that a non-uniformly accelerating observer will measure

$$T = \frac{\pm a(s)/(1 - 2\dot{x}_H)}{2\pi k_b} \tag{5.6}$$

for $a(s) > 0$, $a(s) < 0$ respectively. Here $a(s)$ is the acceleration, s the time, and x_H the location of the Rindler horizon in generalized tortoise coordinates. If the link between thermal time and Unruh temperature holds in this case, we would expect the scaling between thermal and proper time to fluctuate with acceleration.

A third challenge is presented by the case of non-vacuum states. The Radon-Nikodym theorem ensures that the action of the modular automorphism group uniquely determines the generating state. If ϕ, ψ are two (faithful, normal) states on a von Neumann algebra $\mathfrak{M}$, then the associated modular automorphism groups σ^t_ϕ, σ^t_ψ differ by a non-trivial inner automorphism, $\sigma^t_\phi(A) = U\sigma^t_\psi(A)U^*$, for all $A \in \mathfrak{M}$, $t \in \mathbb{R}$. Thus for an eternal, uniformly accelerating observer confined to the Rindler wedge, if the background global state is not ω, then we cannot expect in general that that the wedge modular automorphisms will have a geometric interpretation. Even if they do, they will not be simply related to σ^t_ω by rescaling.

The good news for the TTH is that none of these problems amount to knockdown objections since at present so little is known about the geometric action of modular operators apart from the Bisognano-Wichmann theorem (and its extension to conformal theories). The bad news is that our current ignorance also presents a major challenge to TTH. Cases (1) and (3) in particular pose significant technical questions which require the further development of modular theory to answer. The issues in case (2) are more conceptual in nature. Regarding (1), one intriguing line of investigation concerns the modular automorphism group of double-cone algebras in theories that are asymptotically conformal. As we saw in §4.1, we expect theories with a UV fixed point to be asymptotically conformal, a condition which can be made precise with the help of scaling algebras. In the short distance limit, the double-cone modular operators approximate their wedge counterparts. What does this alignment mean for the double-cone modular structure in the long distance limit?

Regarding case (3), there are clues that states with compact energy (e.g. states satisfying the DHR/BF selection criteria) give rise to well-behaved modular structure on wedges. One of the things that makes the vacuum case so special is that the cyclic, separating vector, Ω, is also the only vector invariant under the global representation of the translations (or more generally the Poincaré group). While this does not hold for any other state, given a vector, Φ, representing a state of compact energy, if l is one of the lightlike vector defining the edge of the wedge, then $U(\lambda l)$, $\lambda \in \mathbb{R}$, is a positively generated subgroup of the translations

mapping $\mathfrak{R}(\mathcal{W})$ onto itself. In addition, $U(\lambda a)\Phi$ is cyclic and separating for $\mathfrak{R}(\mathcal{W})$ and the modular group of $(\mathfrak{R}(\mathcal{W}), U(\lambda l)\Phi)$ can be related to the modular group of $(\mathfrak{R}(\mathcal{W}), \Phi)$ using the Radon-Nikodym derivative.[17] At present, however, nothing more is known.

Given the current results, however, it appears likely that even in the best case scenario, thermal and proper time will not be related by a simple scale factor. To help assuage concerns raised by cases (1)-(3), the defender of the TTH might attempt a more radical move — she could give up on the idea that the thermal time flow determines the metrical properties of time directly. Instead, thermal time would give rise to the order, topological, and group theoretic properties of physical time. The metrical features would be determined by a completely different set of physical relations, such as idealized clock synchronization. This move is supported by Rovelli's comments elsewhere (Rovelli, 2011), which indicate that the primary utility of the TTH is to identify a canonical sense of "flowing" time in a generally covariant theory.[18] Metrical geometry comes second, if applicable at all.

Support for this idea comes from the justification of the clock hypothesis in GR. In most settings the relationship between proper time, τ, and the length of a timelike curve $||\gamma||$ is stipulated. Fletcher (2013) has shown that for any $\epsilon > 0$, there is an idealized lightclock moving along the curve which will measure $||\gamma||$ within ϵ.[19] This justifies the clock

[17] In particular we have

$$\sigma^t_{U(\lambda l)}\Phi(A) = [DU(\lambda l)\Phi : D\Phi]_t \sigma^t_\Phi(A)[DU(\lambda l)\Phi : D\Phi]^*_t, \ \forall A \in \mathfrak{R}(\mathcal{W}).$$

See Borchers (2000) for more details.

[18] Rovelli makes a number of allusions to the concept of an entropy clock discussed by Eddington (1935). Eddington held that the order of temporal events is determined by the thermodynamic arrow of time. An entropy clock measures temporal order by correlating events with decreases in entropy. He describes a simple example:

> An electric circuit is composed of two different metals with their two junctions embedded respectively in a hot and cold body in contact. The circuit contains a galvanometer which constitutes the dial of the entropy-clock. The thermoelectric current in the circuit is proportional to the difference of temperature of the two bodies; so that as the shuffling of energy between them proceeds, the temperature difference decreases and the galvanometer reading continually decreases. (p. 101)

A reliable entropy clock must be in contact with its environment to work properly. In contrast, a reliable metrical clock must be isolated from thermodynamic disturbances. Since the engineering demands pull in separate directions, it might turn out that our phenomenological experience of time is similarly birfurcated.

[19] See also Maudlin (2012, Ch. 5)

hypothesis by linking the metrical properties of spacetime to the readings of tiny light-clocks. The defender of the TTH could argue that the metrical properties of time experienced by localized observers arises via some physical mechanism akin to light clock synchronization. This would explain why the duration of time felt by the observer matches her proper time and not the geometrical flow of thermal time. Moreover, it is consistent with this idea that in the short distance limit, the metrical flow of thermal time agrees with the metrical flow of proper time. As the idealized lightclocks described by Fletcher become smaller and smaller, we approach the regime in which we expect physically well-behaved QFTs to have a UV fixed point. In this limit the double-cone modular operators act geometrically like wedge operators implementing proper time translations along the observer's worldline. It is unlikely that the physics at this scale would impact phenomenology, but the asymptotic connection might turn out to be important for explaining the metrical properties of spacetime (which bigger, more realistic lightclocks measure) as emergent features of some underlying theory of quantum gravity.

It would still be incumbent on the TTH to recover the other, non-metrical features of time. This amounts to showing that the for every physically allowed state, the modular automorphisms act geometric, causally, and order preserving in Trebels's sense. (Note that the defining conditions (i)-(iv) permit conformal transformations, hence "geometric" in here does not refer to metrical geometry.)

As a final comment, we mention one other underdeveloped line of defense which might be useful if the TTH drops its commitment to explaining the metrical features of time. As Rovelli and Connes briefly note, since the modular automorphisms associated with each faithful, normal state on $\mathfrak{M}$ are connected by inner automorphisms, they all project down onto the same 1-parameter group of outer automorphisms on $\mathfrak{M}$. Connes remarks that this captures a sense in which von Neumann algebras are "intrinsically dynamical objects."[20] The TTH might then be revised to the claim that this state-independent canonical flow represents

[20]Connes and Rovelli (1994)

the non-metrical flow of physical time. It is not known, however, under what circumstances the canonical outer flow acts geometric, causal, and order preserving, so it remains to be seen whether or not this is a viable option. The move does have immediate consequences for the global dynamics, however. Since the global algebra is type I, all modular automorphism groups are inner. As a result the canonical group of outer automorphisms determined by the modular structure is trivial. Thus if we identify the outer modular flow as physical time, then at a global level, there is no passage of time. At the local level, the passage of time emerges as a consequence of local ignorance of the global state. For Rovelli, who has long defended timeless mechanics in both classical and quantum contexts, this feature will most likely be seen as a virtue. Less intrepid theorists may demur.

5.3 The Classical Limit

The status of the TTH in the classical limit presents a different kind of challenge. Conceptually, nothing about the basic idea that a coarse-grained statistical state selects a preferred notion of time requires that the theory be quantum mechanical. The proposal for selecting a thermal time parameter using modular theory, however, does appear to rely on the noncommutativity of quantum observables. If we model classical systems using abelian von Neumann algebras, then every state is tracial ($\phi(AB = \phi(BA)$ trivially since $AB = BA$). But a state is tracial iff the corresponding modular automorphism group acts as the identity automorphism, therefore the corresponding flow of thermal time becomes trivial. Does the TTH have a classical counterpart, or do we need quantum mechanics to save time in a generally covariant setting?

Arguing by analogy with standard quantization procedures, Connes and Rovelli suggest that in the classical limit, commutators need to be replaced by poisson brackets. In the quantum case, the thermal time evolution is given by

$$\alpha_t A = e^{it \ln \Delta} A e^{-it \ln \Delta}, \tag{5.7}$$

therefore

$$\frac{d}{dt}\alpha_t A|_{t=0} = i[A, \ln \Delta]. \tag{5.8}$$

In the classical case, we begin with an arbitrary statistical state, ρ, represented by a probability distribution over a classical statespace Γ:

$$\int_\Gamma ds\ \rho(s) = 1, \tag{5.9}$$

where $s \in \Gamma$ is a timeless microstate. By analogy with the Gibbs postulate, we can introduce the "thermal Hamiltonian,"

$$H_\rho = -\ln \rho. \tag{5.10}$$

With respect to the corresponding Hamiltonian vector field, the evolution of an arbitrary classical observable, $f \in C^\infty(\Gamma)$, is given by

$$\frac{d}{dt} f = \{-\ln \rho, f\}, \tag{5.11}$$

and $\rho = \exp(-H_\rho)$. With respect to the poisson bracket structure, the classical algebra of observables is non-abelian. Gallavotti and Pulvirenti (1976) use this non-abelian structure to define an analogue of the KMS condition. Is this connection strong enough to support a version of the TTH in ordinary GR? Or does it only serve to aid us in understanding how the thermal time variable behaves in the transition from quantum theory to classical physics?

The difficulty lies in connecting the thermal time flow for an arbitrary statistical state to our ordinary conception of time. In the quantum case this link was provided by the Bisognano-Wichmann theorem, which does not have a classical analogue. The problem is magnified by the lack of a full understanding of statistical mechanics and thermodynamics in curved spacetime. Rovelli has done some preliminary work on developing a full theory of generally covariant thermodynamics based on the foundation supplied by the TTH, including

an elegant derivation of the Tolman-Ehrenfest effect, but the field is still young.[21] The attention of philosophers could be particularly useful in sorting out the various strategies for extending statistical physics to the generally covariant regime, as well as in analyzing the different abstract features of temporal flows that we might hope to recover, either by a classical TTH, or a quantum TTH in the classical limit.

Setting aside these broader interpretive challenges for now, an important first step lies in obtaining a better understanding the classical thermal time selection procedure outlined by Connes and Rovelli. As it turns out, their commutator-to-Poisson bracket ansatz is on firmer foundational footing than one might initially suppose. In our discussion of the PCT theorem in Ch. 3, we saw that any non-abelian C^*-algebra has the natural structure of a Lie-Jordan algebra. The Jordan product encodes information about the spectra of observables, while the Lie product encodes the generating relation between observables and symmetries. The significance of the commutator, is that it encodes the canonical Lie product. Classical mechanical theories formulated on either a symplectic or Poisson manifold have a natural Lie-Jordan structure as well. The standard product of functions gives the Jordan product, while the Poisson bracket is the Lie product. Together, this structure is called a *Poisson algebra*. As in C^*-algebras, the Jordan product encodes spectral information, while the Lie product describes how classical observables generate Hamiltonian vector fields on statespace. The primary difference between the cases is the associativity/non-associativity of the classical/quantum Jordan products.

These considerations point towards the idea that the appropriate classical analogue of a noncommutative von Neumann algebra, is not a commutative von Neumann algebra, but a Poisson algebra. This is the perspective taken by deformation quantization, one of the most mathematically rigorous approaches to quantization.[22] Additionally, Grgin and Petersen

[21] See Rovelli and Smerlak (2011). The Tolman-Ehrenfest effect is one of the well known modifications of classical thermodynamics expected to arise in GR: the temperature of a closed system in an equilibrium state is not constant, but instead varies with the gradient of the gravitational field. Essentially, a bucket of fluid in thermal equilibrium is hotter at the bottom.

[22] See Landsman (1998)

(1974) show that for any theory in which the observables play the dual role of measurable physical quantities and generators of infinitesimal symmetries, the theory will have a natural Lie-Jordan algebraic formulation. Moreover, assuming that the symmetries of the theory are describable by a central simple Lie algebra, then quantum Lie-Jordan algebras (i.e. C^*-algebras) and classical Lie-Jordan algebras (i.e. Poisson algebras) are effectively the only two mathematical options available.

Initial strides towards a classical analogue of modular theory have been made by Weinstein (1997). Given any smooth density, μ, on a Possion manifold, Γ, Weinstein defines a corresponding *modular vector field* ϕ_μ given by the operator $\phi_\mu : f \rightarrow \mathrm{div}_\mu H_f$ where H_f is the Hamiltonian vector field associated with a classical observable, $f \in C^\infty(\Gamma)$.[23] The anti-symmetry of the Poisson bracket entails that the operator ϕ_μ is a vector field on Γ. Weinstein proposes ϕ_μ as the classical analogue of the modular automorphism group. It characterizes the extent to which the Hamiltonian vector fields are divergence free (with respect to the density μ), vanishing iff all Hamiltonian vector fields are divergence free.

We can connect Weinstein's classical modular theory to the TTH. If Γ is a symplectic manifold and we let μ be the density associated with the canonical Liouville volume form, then $\phi_\mu(f) = 0$ for all observables. This reflects the conservation of energy by Hamiltonian flows in symplectic dynamical systems. Given any statistical state, however, we can define an associated density which leads to a nontrivial modular vector field. For any positive function, h, we have

$$\phi_{h\mu} = \phi_\mu + H_{-\ln h} = H_{-\ln h}. \tag{5.12}$$

Therefore any statistical state, ρ, defines a modular vector field equivalent to the Hamiltonian vector field $H_{-\ln \rho}$ associated with the density $e^{-\ln \rho}\mu$. We immediately recognize $-\ln \rho$ as the thermal Hamiltonian postulated by Connes and Rovelli. Clearly, $e^{it \ln \rho} \rho e^{-it \ln \rho} = \rho$, thus the state is invariant with respect to the flow of $H_{-\ln \rho}$. Additionally, it can be shown that ρ

[23] The divergence of a vector field ξ is defined as $\mathrm{div}_\mu \xi := \pounds_\xi \mu / \mu$ where $\pounds_\xi$ is the Lie derivative with respect to ξ.

satisfies the KMS condition with respect to these dynamics, hence, from the perspective of the associated time flow ρ resembles an invariant equilibrium state just as in the quantum case.

So it turns out that by exploiting the Poisson bracket structure of classical mechanics, we can find a direct analogue for the TTH modular selection procedure. How close is the analogy between the modular vector field of a Poisson manifold and the modular automorphism group of a von Neumann algebra? Can other remnants of modular structure be found within classical mechanics? Weinstein's proto-theory deserves further study by philosophers of physics, not only because of its connection to the TTH, but because of insight it might provide into the physical interpretation of modular theory in QFT.

5.4 Conceptual Challenges

As we have seen in the previous two sections, the TTH faces a number of technical challenges (some of which look easier to overcome than others). There are, however, several deeper conceptual problems looming in the background which pose a more serious challenge to the viability of the hypothesis. Here, we will discuss three of the most pressing.

The first, which we'll call the *generality problem*, draws upon our previous discussion of the TTH in the classical limit. While mathematically speaking, Weinstein's modular vector field gives us a method for selecting a canonical thermal time flow in a classical theory, physical speaking, there is no reason why we should view the corresponding thermal time as physical time. As we have seen, any statistical state determines thermal dynamics according to which it is a KMS state, however, if ρ is a non-equilibrium state, the resultant thermal time flow does not align with our ordinary conception of time. By the lights of thermal time, a cube of ice in a cup of hot coffee is an invariant equilibrium state! The same problem arises in the quantum domain — only for states which are true equilibrium states will the thermal time correspond to physical time. Of course without an intrinsic characterization

of equilibrium, independent of an antecedent notion of time, the TTH cannot appeal to this fact. Connes and Rovelli's original attempt to circumvent the problem using modular theory winds up severely overgeneralizing. The resulting concept of thermal time is simply too broad to be identified with physical time. Only in restricted physical contexts will thermal time and physical time be equivalent.

It appears inevitable that the TTH will have to be tempered. Rather than letting any state determine a corresponding flow of thermal time, only certain reference states should be permitted. The most obvious candidates are equilibrium states. Apart from the problem of providing an intrinsic characterization of such states, if a system is not in equilibrium, it is hard to envision how a counterfactual state of affairs could determine the actual flow of time. This might provide more reasons for the defender of the TTH to explore the state independent, outer modular flow. Alternatively, he could try to argue that local non-equilibrium behavior can be viewed as small fluctuations in some background state. On this approach, the local flow of time in my office according to which the ice melts and the coffee cools isn't defined by the thermal state of the ice/coffee system, but the thermal state of some larger enveloping system (the entire universe perhaps). Hints in this direction can be found in Rovelli (1993). In this paper, Rovelli argues that in a Friedman-Robertson-Walker universe, the thermal time induced by the equilibrium state of the cosmic microwave background will be equivalent to the FRW time.[24] While the connection is intriguing, it seems unlikely that an explanation of this sort will be able to account for the flow of time experienced by localized, mortal observers like us. It would be truly remarkable to discover that our faculties of perception are sensitive to the thermal features of the CMB.

[24]The FRW metric,

$$ds^2 = dt^2 + a^2(t)[dr^2 + S_k^2(r)d\Omega^2]$$

(where $S_k(r)$ varies depending on the sign of the curvature constant, k) describes an expanding homegenous, isotropic universe with dimensionless scale-factor a. The proper time experienced by free-falling observers co-moving with the expansion aligns with the coordinate time t. This cosmological time parameter is known as *FRW time*.

A closely related worry has been voiced by Earman (2011) and Ruetsche (2014). In the physical situations where we can justify the modular automorphism group as a kind of dynamics, it seems this is only possible because we already have a sufficiently rich spatiotemporal structure in the background. This casts doubt on whether the TTH can provide a coherent definition of time in situations where such structure is absent (as required to solve the full problem of time). In the Bisognano-Wichman scenario, we immediately recognize the geometric significance of the modular automorphism group because its flow is everywhere timelike. The orbits of σ_t correspond to a clear class of observer worldlines and $d\tau/dt$ is constant along those worldine, yielding a simple scaling relation between t and τ. In other cases, even when the modular operators act geometrically, it can be hard to recognize σ_t as dynamical. The conformal examples from §5.2 show that the scaling relation between thermal and proper time can be highly non-trivial. An even stranger case is presented by the modular group associated with the forward lightcone in the vacuum sector of a free, massless scalar theory. Since the field propagates at the speed of light, the theory possesses timelike commutativity in addition to spacelike commutativity. By the Reeh-Schlieder theorem, the restriction of the vacuum state to $\mathfrak{R}(V^+)$ is therefore cyclic and separating. The corresponding modular group acts as dilations $x^\mu \mapsto e^{-2\pi t}x^\mu$ on vectors in V^+.[25] If we interpret the orbits of σ_t as the worldlines of inertial observers departing from the origin of the lightcone, the relationship between thermal time and proper time has the form $t(\tau) = -\ln \tau/2\pi$.

Although this interpretation is available, it appears strained. Moreover, extracting it requires antecedent knowledge of background causal structure. In generally covariant spacetimes with no global timelike killing fields, no global isometries, etc. such an interpretation may no longer be possible. The problem is exacerbated if the TTH is modified in response to the generality problem. Unless the modular automorphism group can always be viewed dynamically, the defender of the TTH will be hard pressed to find constraints capable of separating the dynamical cases from the non-dynamical cases which are independent of all

[25] Haag (1996, Thm. 4.2.2).

background temporal structure. We'll call this second problem, the *background-dependence problem*.

The third and final problem is the *gauge problem*. In spite of all the challenges discussed above, the TTH does succeed in providing a means to select a privileged 1-parameter flow on the space of full, gauge invariant observables of a generally covariant theory. What makes this flow interpretable as a *dynamical* flow, however, is its description as a sequence of correlations between partial observables (namely, between the thermal time parameter and other partial observables appearing in the theory). The difficulty is that these partial observables are not diffeomorphism invariant. Assuming that we treat diffeomorphisms in generally covariant theories as standard gauge symmetries (which is how we got into the problem of time in the first place), then the partial observables are just descriptive fluff. They do not directly represent physical features of our world.

The problem is *not* the resultant timelessness of fundamental physics. The TTH adopts this dramatic conclusion willingly. The problem is that the TTH is supposed to explain how the appearance of time and change emerge from timeless foundations. But the explanation we've been given is couched in gauge-dependent language, and it is not apparent how we can extract a gauge invariant story from it.

An analogy with classical spacetime physics will serve to illuminate the central issue. It is widely thought that the invariance of Newton's second law with respect to Galilean boosts indicates that, fundamentally, there are no facts about absolute velocity in classical spacetime. There are however, gauge-invariant facts about relative velocities. By selecting a preferred reference frame defining a state of rest (i.e. fixing a gauge), we can introduce absolute velocities into our theoretical description of the world. We can even go on to use these absolute velocities to compute gauge-invariant relative velocities. But we cannot use correlations between absolute veleocities to explain facts about relative velocity. (At least if we take explanation in a suitably robust sense.) The fact that the relative velocity between Alice and Bob is 35 kph northwest cannot be explained by citing the absolute velocities of

each party and taking the vector difference. There are no facts about absolute velocity to appeal to. Instead, our explanation must employ gauge invariant features of the metrical structure of Galilean spacetime.

In a generally covariant setting we can introduce partial observables and use correlations between them to calculate and predict emergent dynamical behavior, but we cannot use these correlations to explain that behavior. At this point we still lack a gauge invariant picture of generally covariant theories akin to the one provided by Galilean spacetime in the example above. The TTH, at least in its present form, does not provide one.

There is one potentially serious disanalogy with the classical velocity example. In the velocity case, the explananda are gauge-invariant facts about relative velocity. The facts in question in the thermal time case are perspectival facts about how the flow of time *appears to us*. Rovelli emphasizes that although the value of partial observables cannot be predicted given the state, they can still be measured in some sense, opening the possibility that they might be explanatorily significant via their connection to measurement processes. How this is possible is unclear. When we witness an object change position, we measure two gauge invariant quantities $q(t_1)$ and $q(t_2)$. We can describe these as measurements of correlations between position and time partial observables, but to do so requires a gauge-dependent deparametrization of the constraint Hamiltonian. If at the end of the day the partial observables are simply descriptive fluff, how can a measuring device be sensitive to them? (No Newtonian apparatus, however ingeniously designed, can measure absolute velocities.) It might be possible to view the partial observables in question as intrinsically perspectival (e.g. velocity-relative-to-*my*-reference-frame). While this would serve to connect them to observation directly, it would effectively build the explanandum into the explanans.[26]

Can a revised form of the TTH provide us with the explanatory tools to understand the flow of thermal time without reference to partial observables, or does the entire framework

[26]Moreover, an explanation of velocity-relative-to-my-reference-frame should appeal to gauge invariant facts about relative velocities and different reference frames along with a single indexical fact picking out which frame I'm in. But the explanatory role of the thermal time partial observable extends well beyond this minimal indexicality.

of timeless mechanics require us to revise our conception of how ontology, explanation, and gauge symmetries are related?[27] Whether or not quantum thermodynamics can save time may ultimately rest on the solutions to these new reincarnations of vexingly familiar philosophical problems.

5.5 Modular Dynamics

In light of the generality, background-dependence, and gauge problems, it seems doubtful that the TTH will be able to succeed in providing a comprehensive solution to the problem of time. Within the context of QFT on a fixed background spacetime, however, the core intuition that the modular automorphism group encodes information about the local dynamics is intriguing. As we've already seen, in theories satisfying nuclearity, the modular spectrum provides a bound on the local energy. The geometric significance of the modular operators expressed by the Bisognano-Wichmann relations offer another clue. While these facts are encouraging, the exact relationship between Δ^{it} and the local Hamiltonian remains elusive.

A minimal necessary condition on any concept of localized dynamics is convergence to the global dynamics in the infinite volume limit. The global dynamics are represented by a unitary representation of the translations, $U(a)$, on $\mathcal{H}$. Given a continuous family of nested double-cones of increasing size (centered at x), and a family of unitary groups, $\{U_R(t)\}$, indexed by radius R, we require that

$$U_R(t)A\Omega \to U(tx)A\Omega, \tag{5.13}$$

[27] Rovelli's most recent work appears to be drifting in the latter direction. Rovelli (2013) suggests that gauge-dependent quantities are more than just mathematical redundancies, arguing that they are critical for understanding interactions between physical systems — "they describe handles though which systems couple: they represent real relational structures to which the experimentalist has access in measurement by supplying one of the relata in the measurement procedure itself." It should be emphasized that this marks a radical break from the received view on gauge. For instance, absolute velocity would count as an experimentally accessible handle by the lights of Rovelli's new account.

in the Hilbert space topology, as $R \to \infty$, for every $A \in \mathfrak{R}(\mathcal{O})$, $t \in \mathbb{R}$. For conformal field theories, Borchers (2000) shows that a suitable rescaling of the modular unitaries obeys (5.13):

$$U_R(t) := \Delta_R{}^{-i\frac{t}{\pi R}} \tag{5.14}$$

In nonconformal theories, we can replace the nested family of double-cones with a family of nested wedges. (The parameter R now represents the increasing spacelike distance between the vertices of the nested wedges.) We find a similar convergence result but with a slightly different rescaling:

$$U_R(t) := \Delta_R{}^{-i\frac{t}{2\pi R}} \tag{5.15}$$

(Why an extra factor of 2 appears in this case is unclear.) Together, these results add weight to the intuition that the physical interpretation of the modular operators is linked to dynamics.

The significance of a potential connection lies in the localization problem discussed in chapter 4. Since the dynamics are globally defined, it is difficult to recover a local description of them within the setting of AQFT. As we've seen, in theories satisfying the split/nuclearity condition, the universal localization map can be used to find local generators in a slightly larger enveloping region. In particular, we have an expression for the local Hamiltonian associated with a given standard split inclusion, $H_\Lambda = \psi_\Lambda H$. While this solves the problem of how we can measure the energy in a localized region, it appears that in order to have any detailed knowledge of the local dynamics, we must already know what the global dynamics

are.[28] Even if modular automorphisms do not implement the local dynamics directly, any information or constraints they provide would be useful in countering this objection.

Thus far, all attention has been focused on the question, "do the modular objects associated with $\mathfrak{R}(\mathcal{O})$ generate a local dynamics for $\mathcal{O}$?" Given our discussion of the nuclearity/split condition in Ch. 4, a more natural question would be, "given a standard split inclusion Λ, do the modular objects associated with either $\mathfrak{R}(\tilde{\mathcal{O}})$ or $\mathfrak{N}_\Lambda$ generate a local dynamics on $\mathcal{O}$?" In the conformal case, we saw that the double-cone modular unitaries approximated proper time translations for most of the mortal observer's life before diverging at the end. This suggests that on any suitably smaller subdiamond, these unitaries generate a local representation of the translations, just as the local Hamiltonian H_Λ does. More work is needed in this direction. Even if the TTH cannot solve the problem of time, it might yet provide the key for a rigorous understanding of the localization of dynamics in standard QFT.

[28] The same problem arises for the local charge structure — we can only know what the local superselection rules are if we know the global rules. Buchholz et al. (1986) show, however, that this problem reduces to the knowledge of the local Hamiltonian. Consider the von Neumann algebra $\mathfrak{M}$ generated by the union of the local Hamiltonian H_Λ and the local C^*-algebra $\mathfrak{A}(\mathcal{O})$. Since ψ_Λ acts trivially on $\mathfrak{A}(\mathcal{O})$ and is ultraweakly continuous,

$$\psi_\Lambda(\bigvee_t U(t)\mathfrak{A}(\mathcal{O})U(t)^{-1}) = \bigvee_t U_\Lambda(t)\mathfrak{A}(\mathcal{O})U_\Lambda(t)^{-1} \subset \mathfrak{M}.$$

By (strong) additivity and the spectrum condition, $\bigvee_t U(t)\mathfrak{A}(\mathcal{O})U(t)^{-1} = \mathfrak{A}''$. But since the translations commute with the gauge group, $U(t) \in \mathfrak{A}''$, and so $\mathfrak{M} \subset \psi_\Lambda(\mathfrak{A}'')$. Thus $\mathfrak{M}$ is isomorphic to the localization of the global algebra.

Bibliography

Alfsen, E. and Shultz, F. (2001). *State Spaces of Operator Algebras*. Birkhäuser.

Alfsen, E. and Shultz, F. (2003). *Geometry of State Spaces of Operator Algebras*. Birkhäuser.

Araki, H. (1964a). On the algebra of all local observables. *Progress of Theoretical Physics*, 32(5):844–854.

Araki, H. (1964b). Type of von Neuman algebra associated with free field. *Progress of Theoretical Physics*, 32(6):956–965.

Arntzenius (2011). The CPT theorem. In Callender, C., editor, *The Oxford Handbook of Philosophy of Time*. Oxford University Press.

Arntzenius, F. and Greaves, H. (2009). Time reversal in classical electromagnetism. *British Journal for the Philosophy of Science*, 60(3):557–584.

Bachmann, F., Pejas, W., Wolff, H., and Bauer, A. (1986). *Fundamentals of Mathematics: Geometry*, volume 2, chapter 5. MIT Press.

Bain, J. (2013). CPT invariance, the spin-statistics connection, and the ontology of relativistic quantum field theories. *Erkenntnis*, 78(4):797–821.

Baker, D. J. (2014). Can we make do without observables? evaluating spacetime state realism. In preparation.

Baker, D. J. and Halvorson, H. (2010). Antimatter. *British Journal for the Philosophy of Science*, 61:93–121.

Baker, D. J., Halvorson, H., and Swanson, N. W. (2014). On the conventionality of parastatistics. forthcoming in *British Journal for the Philosophy of Science*.

Barrett, J. A. (2002). On the nature of measurement records in relativistic quantum field theory. In *Ontological Aspects of Quantum Field Theory*. World Scientific.

Bell, J. (1955). Time reversal in field theory. *Proclamations of the Royal Society London, A*, 231(1187):479–495.

Belot, G. (2005). The representation of time and change in mechanics. In Earman, J. and Butterfield, J., editors, *Handbook of the Philosophy of Physics*. North-Holland.

Bisognano, J. and Wichmann, E. (1975). On the duality condition for a hermitian scalar field. *Journal of Mathematical Physics*, 16(4):985–1007.

Bisognano, J. and Wichmann, E. (1976). On the duality condition for quantum fields. *Journal of Mathematical Physics*, 17(303-321).

Blackadar, B. (2006). *Operator Algebras*. Springer.

Borchers, H. J. (1967). A remark on a theorem of B. Misra. *Communications in Mathematical Physics*, 4:315–323.

Borchers, H. J. (1992). The CPT-theorem in two-dimensional theories of local observables. *Communications in Mathematical Physics*, 143:315–332.

Borchers, H. J. (1996). *Translation Group and Particle Representations in Quantum Field Theory*. Springer-Verlag.

Borchers, H. J. (2000). On revolutionizing quantum field theory with Tomita's modular theory. *Journal of Mathematical Physics*, 41(6):3604–3673.

Borchers, H. J. and Wollenberg, M. (1991). On the relation between types of local algebras in different representations. *Communications in Mathematical Physics*, 137:161–173.

Borchers, H. J. and Yngvason, J. (1992). From quantum fields to local von Neumann algebras. *Reviews in Mathematical Physics*, 4(1):15–47.

Borchers, H. J. and Yngvason, J. (1994). Transitivity of locality and duality in quantum field theory. Some modular aspects. *Reviews in Mathematical Physics*, 6:597.

Borchers, H. J. and Yngvason, J. (2000). On the PCT-theorem in the theory of local observables. http://arxiv.org/abs/math-ph/0012020.

Bratteli, O. and Robinson, D. W. (1981). *Operator Algebras and Quantum Statistical Mechanics*. Springer-Verlag.

Brunetti, R. and Fredenhagen, K. (2004). The algebraic approach to quantum field theory. http://arxiv.org/abs/math-ph/0411072.

Brunetti, R., Fredenhagen, K., and Verch, R. (2003). The generally covariant locality principle – a new paradigm for local quantum field theory. *Communications in Mathematical Physics*, 237:31–68.

Brunetti, R., Guido, D., and Longo, R. (1993). Modular structure and duality in conformal quantum field theory. *Communications in Mathematical Physics*, 156:201–219.

Buchholz, D., D'Antoni, C., and Longo, R. (1990). Nuclear maps and modular structure II. *Communications in Mathematical Physics*, 129:115.

Buchholz, D., Doplicher, S., and Longo, R. (1986). On Noether's theorem in quantum field theory. *Annals of Physics*, 170:1–17.

Buchholz, D., Doplicher, S., and Roberts, J. (2013). New light on infrared problems: Sectors, statistics, symmetries and spectrum. page http://arxiv.org/abs/1304.2794.

Buchholz, D., Dreyer, O., Florig, M., and Summers, S. J. (2000). Geometric modular action and spacetime symmetries. *Reviews in Mathematical Physics*, 12:475–560.

Buchholz, D. and Fredenhagen, K. (1982). Locality and the structure of particle states. *Communications in Mathematical Physics*, 84(1):1–54.

Buchholz, D. and Summers, S. J. (1993). An algebraic characterization of vacuum states in minkowski space. *Communications in Mathematical Physics*, 155:449–458.

Buchholz, D. and Summers, S. J. (2008). Warped convolutions: A novel tool in the construction of quantum field theories. In Seiler, E. and Sibold, K., editors, *Quantum Field Theory and Beyond*, pages 107–121. World Scientific, Singapore.

Buchholz, D. and Verch, R. (1995). Scaling algebras and renormalization group in algebraic quantum field theory. *Reviews in Mathematical Physics*, 7:1195.

Buchholz, D. and Wichmann, E. H. (1986). Causal independence and the energy-level density of states in local quantum theory. *Communications in Mathematical Physics*, 106:321.

Buchholz, D. and Yngvason, J. (1994). There are no causality problems in Fermi's two atom system. *Physical Review Letters*, 72:596–599.

Clifton, R. (2000). The modal interpretation of algebraic quantum field theory. *Physics Letters A*, 271:167–177.

Clifton, R., Bub, J., and Halvorson, H. (2003). Characterizing quantum theory in terms of information-theoretic constraints. *Foundations of Physics*, 33:1561–1591.

Clifton, R. and Halvorson, H. (2001). Entanglement and open systems in quantum field theory. *Studies in History and Philosophy of Modern Physics*, 32:1–31.

Coleman, S. and Mandula, J. (1967). All possible symmetries of the s-matrix. *Physics Review*, 159(5):1251–1256.

Connes, A. (1973). Un classification de factors de type III. *Annales Scientifiques de l'Ecole Normale Supérieure*, 6:133–252.

Connes, A. (1974). Caractérisation des algèbres de von Neumann commes espaces vectoriels ordonnés. *Annales de l'Institute Fourier*, 24:121–155.

Connes, A. and Rovelli, C. (1994). Von Neumann algebra automorphisms and time-thermodynamics relation in generally covariant quantum theories. *Classical and Quantum Gravity*, 11(12):2899.

Connes, A. and Størmer, E. (1978). Homogeneity of the state space of factors of type III_1. *Journal of Functional Analysis*, 28(2):187–196.

Conway, J. (1990). *A Course in Functional Analysis*. Springer.

D'Antoni, C., Doplicher, S., Fredenhagen, K., and Longo, R. (1987). Convergence of local charges and continuity properties of w^*-inclusions. *Communications in Mathematical Physics*, 110:325–348.

Davidson, D. R. (1995). Modular covariance and the algebraic PCT/spin-statistics theorem. http://arxiv.org/abs/hep-th/9511216.

Doplicher, S. (1982). Local aspects of superselection rules. *Communications in Mathematical Physics*, 85:73.

Doplicher, S. and Longo, R. (1983). Local aspects of superselection rules II. *Communications in Mathematical Physics*, 88:399.

Doplicher, S. and Longo, R. (1984). Standard and split inclusions of von Neumann algebras. *Inventiones Mathematicae*, 75:493–536.

Doplicher, S. and Roberts, J. E. (1990). Why there is a field algebra with a compact gauge group describing the superselection structure in particle physics. *Communications in Mathematical Physics*, 131:51–107.

Driessler, W. (1977). On the type of local algebras in quantum field theory. *Communications in Mathematical Physics*, 53:295–297.

Earman, J. (2002). Thoroughly modern McTaggart. *Philosopher's Imprint*, 2. http://www.philosophersimprint.org/002003/.

Earman, J. (2011). The Unruh effect for philosophers. *Studies in History and Philosophy of Modern Physics*, 42:81–97.

Earman, J. and Fraser, D. (2006). Haag's theorem and its implications for the foundations of qauntum field theory. *Erkenntnis*, 64:305–344.

Eddington, A. (1935). *The Nature of the Physical World*. MacMilan.

Ellers, E. W. (2004). Conjugacy classes of involutions in the Lorentz group $\Omega(V)$ and in $SO(V)$. *Linear Algebra and its Applications*, 383:77–83.

Fernandez, R., J. F. and Sokal, A. (1992). *Random Walks, Critical Phenomena, and Triviality in Quantum Field Theory*. Springer-Verlag, New York.

Feynman, R. (1949). The theory of positrons. *Physical Review*, 76(6):749–759.

Fletcher, S. (2013). Light clocks and the clock hypothesis. *Foundations of Physics*, 43:1369–1383.

Folland, G. B. (1995). *A course in abstract harmonic analysis*. CRC Press, Boca Raton, FL.

Fraser, D. (2009). Quantum field theory: Underdetermination, inconsistency and idealization. *Philosophy of Science*, 76:536–567.

Fredenhagen, K. (1985a). On the modular structure of local algebras of observables. *Communications in Mathematical Physics*, 97:79–89.

Fredenhagen, K. (1985b). A remark on the cluster theorem. *Communications in Mathematical Physics*, 97:461–463.

Fredenhagen, K. (2004). Locally covariant quantum field theory.

Gallavotti, G. and Pulvirenti, M. (1976). Classical KMS condition and Tomita-Takesaki theory. *Communications in Mathematical Physics*, 46:1–9.

Giannitrapani (1998). Quantum coordinates of an event in local quantum physics. *Journal of Mathematical Physics*, 39:5180.

Greaves, H. (2008). *Spacetime Symmetries and the CPT Theorem*. PhD thesis, Rutgers University.

Greaves, H. (2010). Towards a geometrical understanding of the CPT theorem. *British Journal for the Philosophy of Science*, 61(1):27–50.

Greaves, H. and Thomas, T. (2012). The CPT theorem. http://arXiv:1204.4674.

Greenberg, O. (2002). CPT violation implies violation of Lorentz invariance. *Physical Review Letters*, 89(23):231602.

Grgin, E. and Petersen, A. (1974). Duality of observables and generators in classical and quantum mechanics. *Journal of Mathematical Physics*, 15:764.

Guenin, M. and Misra, B. (1963). On the von Neumann algebras generated by field operators. *Il Nuovo Cimento*, 30(5):1272–1290.

Guido, D. and Longo, R. (1992). Relativistic invariance and charge conjugation in quantum field theory. *Communications in Mathematical Physics*, 148(3):521–551.

Guido, D. and Longo, R. (1994). An algebraic spin statistics theorem. *Communications in Mathematical Physics*, 172:517–533.

Guido, D. and Longo, R. (2002). Modular localization and Wigner particles. *Reviews in Mathematical Physics*, 14:759–785.

Haag, R. (1996). *Local Quantum Physics*. Springer-Verlag, Berlin.

Haag, R. and Swieca, J. (1965). When does a quantum field theory describe particles? *Communications in Mathematical Physics*, 1:308.

Haagerup, U. (1987). Connes' bicentralizer problem and uniqueness of the injective factor of type III_1. *Acta Mathematica*, 158:95.

Halvorson, H. and Clifton, R. (2002). No place for particles in relativistic quantum theories? *Philosophy of Science*, 69.

Halvorson, H. and Müger, M. (2006). Algebraic quantum field theory. In Butterfield, J. and Earman, J., editors, *Philosophy of Physics*, pages 731–922. Elsevier.

Hanneke, D., Fogwell Hoogerheide, S., and Gabrielse, G. (2011). Cavity control of a single-electron quantum cyclotron: Measuring the electron magnetic moment. *Physical Review A*, 83:052122.

Hegerfeldt, G. (1994). Causality problems for Fermi's two-atom system. *Physical Review Letters*, 72(5):596–599.

Hislop, P. D. and Longo, R. (1982). Modular structure of the local algebras associated with a free massless scalar field theory. *Communications in Mathematical Physics*, 84:71.

Horuzhy, S. S. (1990). *Introduction to Algebraic Quantum Field Theory*. Kluwer Academic Publishers.

Jian-yang, Z., Aidong, B., and Zheng, Z. (1995). Rindler effect for a nonuniformly accelerating observer. *International Journal of Theoretical Physics*, 34:2049–2059.

Jost, R. (1965). *The General Theory of Quantized Fields*. AMS.

Kadison, R. V. (1963). Remarks on the type of von Neumann algebras of local observables in quantum field theory. *Journal of Mathematical Physics*, 4(12):1511–1516.

Kadison, R. V. and Ringrose, J. R. (1997a). *Fundamentals of the theory of operator algebras. Vol. I.* American Mathematical Society, Providence, RI.

Kadison, R. V. and Ringrose, J. R. (1997b). *Fundamentals of the theory of operator algebras. Vol. II.* American Mathematical Society, Providence, RI.

Knight, J. (1961). Strict localization in quantum field theory. *Journal of Mathematical Physics*, 2(469):459–471.

Kraus, K. (1964). General quantum field theory and strict locality. *Zeitschrift für Physik 181*, 181:1–12.

Kuckert, B. (1997). Borchers' commutation relations and modular symmetries in quantum field theory. *Letters in Mathematical Physics*, 41:307–320.

Landsman, N. (1998). *Mathematical Topics Between Classical and Quantum Mechanics*. Springer-Verlag, New York.

Lechner, G. (2008). Construction of quantum field theories with factorizing S-matrices. *Communications in Mathematical Physics*, 277(3):821–860.

Lévy-Leblond, J.-M. (1967). Galilean quantum field theories and a ghostless Lee model. *Communications in Mathematical Physics*, 4:157–176.

Longo, R. (1979). Notes on algebraic invariants for non-commutative dynamical systems. *Communications in Mathematical Physics*, 69:195–207.

Longo, R. (1987). The joint modular structure of and inclusion of von Neumann algebras. *Contemporary Mathematics*, 62:529–538.

Malament, D. (1977). Causal theories of time and the conventionality of simultaneity. *Noûs*, 11(3):293–300.

Malament, D. (2012). *Topics in the Foundations of General Relativity and Newtonian Gravitation Theory*. Univeristy of Chicago Press.

Malament, D. B. (1996). In defense of dogma: Why there cannot be a relativistic quantum mechanics of (localizable) particles. In Clifton, R., editor, *Perspectives on Quantum Reality*, pages 1–10. Kluwer Academic Publishers, Dordrecht.

Martinetti, P. (2004). A brief remark on Unruh effect and causality.

Martinetti, P. and Rovelli, C. (2003). Diamond's temperature: Unruh effect for bounded trajectories and thermal time hypothesis. *Classical and Quantum Gravity*, 20(22):4919.

Maudlin, T. (2012). *Philosophy of Physics: Space and Time*. Princeton University Press.

Mund, J. (2001). The Bisognano-Wichmann theorem for massive theories. *Annales Henri Poincaré*, 2:907–926.

Murray, F. J. and von Neumann, J. (1936). On rings of operators. *Annals of Mathematics*, 37(116-229).

Paetz, T.-T. (2010). An analysis of the 'thermal-time concept' of Connes and Rovelli. Master's thesis, Georg-August-Universität Göttingen.

Peskin, M. and Schroeder, D. (1995). *An Introduction to Quantum Field Theory*. Westview Press.

Rieffel, M. A. and van Dale, A. (1977). A bounded operator approach to Tomita-Takesaki theory. *Pacific Journal of Mathematics*, 69(1):187–221.

Roberts, B. (2013). Three myths about time reversal invariance.

Roberts, J. E. (1982). Localization in algebraic field theory. *Communications in Mathematical Physics*, 85:87–98.

Roberts, J. E. and Roepstorff, G. (1969). Some basic concepts of algebraic quantum theory. *Communications in Mathematical Physics*, 11:321–338.

Rovelli, C. (1993). The statistical state of the universe. *Class.Quant.Grav.*, 10:1567.

Rovelli, C. (2011). Forget time: Essay written for the FQXi contest on the nature of time. *Foundations of Physics*.

Rovelli, C. (2013). Why gauge? http://arxiv.org/abs/1308.5599.

Rovelli, C. and Smerlak, M. (2011). Thermal time and Tolman–Ehrenfest effect: 'temperature as the speed of time'. *Classical and Quantum Gravity*, 28(7):075007.

Ruetsche, L. (2011a). *Interpreting Quantum Theories*. Oxford UP, Oxford.

Ruetsche, L. (2011b). Why be normal? *Studies in History and Philosophy of Modern Physics*, (42):107–115.

Ruetsche, L. (2014). Warming up to thermal the thermal time hypothesis. Quantum Time Conference, University of Pittsburgh, March 28-29.

Saffary, T. (2005). *Modular Action on the Massive Algebra*. PhD thesis, Hamburg.

Schroer, B. and Weisbrock, H.-W. (2000). Modular constructions of quantum field theories with interactions. *Reviews in Mathematical Physics*, 12:301–326.

Sewell, G. (1980). Relativity of temperature and the Hawking effect. *Physics Letters A*, 79:23–24.

Skau, C. F. (1980). Geometric aspects of the Tomita-Takesaki theory I. *Mathematica Scandinavica*, 47:311–328.

Streater, R. F. (1967). Local fields with the wrong connection between spin and statistics. *Communications in Mathematical Physics*, 5:88–96.

Streater, R. F. and Wightman, A. S. (1989). *PCT, Spin and Statistics, and All That*. Addison-Wesley, New York.

Strocchi, F. (1993). *Selected Topics on the General Properties of Quantum Field Theories*. Lecture Notes in Physics, Vol. 51. World Scientific.

Strocchi, F. (2013). *An Introduction to Non-Perturbative Foundations of Quantum Field Theory*. Oxford Science Publications.

Summers, S. J. (1990). Independence of local aglebras in quantum field theory. *Reviews in Mathematical Physics*, 2:201–247.

Summers, S. J. (2006). Tomita-Takesaki modular theory. In *Encyclopedia of Mathematical Physics*. Academic Press/Elsevier Science.

Summers, S. J. (2012a). A perspective on constructive quantum field theory. http://arxiv.org/abs/1203.3991.

Summers, S. J. (2012b). Yet more ado about nothing: The remarkable relativistic vacuum state. In Halvorson, H., editor, *Deep Beauty: Understanding the Quantum World*, pages 317–342. Cambridge University Press.

Summers, S. J. and White, R. K. (2003). On deriving space-time from quantum observables and states. *Communications in Mathematical Physics*, 237:203–220.

Swanson, N. W. and Halvorson, H. (2012). On North's 'the structure of physics'. http://philsci-archive.pitt.edu/9314/.

Takesaki, M. (1970). *Tomita's Theory of Modualr Hilbert Algebras and its Applications*, volume 128 of *Lecture Notes in Mathematics*. Springer.

Takesaki, M. (2000). *Theory of Operator Algebras, Vol. II*. Springer-Verlag.

Trebels, S. (1997). *Über die Geometrische Wirkung Modularer Automoprhismen*. PhD thesis, Göttingen.

Unruh, W. G. (1976). Notes on black hole evaporation. *Physical Review D*, 14:870.

Wallace, D. (2006). In defence of naiveté: On the conceptual status of Lagrangian quantum field theory. *Synthese*, 151:33–80.

Wallace, D. (2009). QFT, antimatter, and symmetry. *Studies in History and Philosophy of Modern Physics*, 40:209–242.

Wallace, D. (2011). Taking particle physics seriously: A critique of the algebraic approach to quantum field theory. *Studies in History and Philosophy of Modern Physics*, 42:116–125.

Wallace, D. (2012). *The Emergent Multiverse*. Oxford University Press.

Wallace, D. and Timpson, C. G. (2010). Quantum mechanics on spacetime I: Spacetime state realism. *British Journal for the Philosophy of Science*, 61:697–727.

Weatherall, J. O. (2011). The motion of a body in Newtonian theories. *Journal of Mathematical Physics*, 52(3):32502–32518.

Weinberg, S. (2005). *The Quantum Theory of Fields*, volume I. Cambridge University Press.

Weinstein, A. (1997). The modular automorphism group of a Poisson manifold. *Journal of Geometry and Physics*, 23:379–394.

Weisbrock, H.-W. (1992). A comment on recent work of Borchers. *Letters in Mathematical Physics*, 25:157–159.

Yngvason, J. (1994). A note on essential duality. *Letters in Mathematical Physics*, 31:127–141.

Yngvason, J. (2005). The role of type III factors in quantum field theory. *Reports in Mathematical Physics*, 55(1):135–137.