Can Quantum Thermodynamics Save Time?

Noel Swanson*

March 31, 2020

Abstract

The *thermal time hypothesis (TTH)* is a proposed solution to the problem of time: a coarse-grained state determines a thermal dynamics according to which it is in equilibrium, and this defines the flow of time in generally covariant quantum theories. This paper raises a series of objections to the TTH as developed by Connes and Rovelli (1994). Two technical challenges concern the relationship between thermal time and proper time and the possibility of implementing the TTH in classical theories. Three conceptual problems concern the flow of time in non-equilibrium states and the extent to which the TTH is background-independent and gauge-invariant.

1 Introduction

In both classical and quantum theories defined on fixed background spacetimes, 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 state space (Schrödinger picture) or a parametrized group of automorphisms of the algebra of observables (Heisenberg picture). In generally covariant theories, where diffeomorphisms of the underlying spacetime manifold are treated as gauge symmetries, this picture breaks down. There is

*Department of Philosophy, University of Delaware, 24 Kent Way, Newark, DE 19716, USA, nswanson@udel.edu

no longer a canonical time-translation subgroup at the global level, nor is there a gauge-invariant way to represent dynamics locally in terms of the Schrödinger or Heisenberg pictures. Without a preferred flow on the space of states representing time, the standard way to represent physical change via functions on this space taking different values at different times also fails. This is the infamous *problem of time*.[1]

Connes and Rovelli (1994) propose a radical solution to the problem: the flow of time (not just its direction) has a thermodynamic origin. Any coarse-grained, statistical state naturally defines a notion of time according to which it is in equilibrium. The *thermal time hypothesis (TTH)* identifies this state-dependent thermal time with physical time. Drawing upon tools from Tomita-Takesaki modular theory, Connes and Rovelli demonstrate how the TTH can be rigorously implemented in generally covariant quantum theories.

The idea is an intriguing one that, to date, has received little attention from philosophers. Not only does the TTH represent a striking conjecture about the origins of time, it also supplies tantalizing clues about the physical significance of Tomita-Takesaki modular theory. One of the most powerful mathematical tools we have to study the structure of operator algebras used in quantum theory, modular theory has found an increasingly diverse array of physical applications.[2] Despite its importance, the basic physical ideas behind modular theory remain murky. If the TTH is right, then modular automorphism groups are employed by generally covariant quantum theories to characterize emergent dynamics.

This paper represents a modest initial attempt to sally forth into rich philosophical territory. Its goal is to voice a number of technical and conceptual challenges faced by the TTH and to propose some strategies to respond. In §2, I provide a self-contained introduction to the TTH, emphasizing the connection between Connes and Rovelli's original proposal and Rovelli's later work on timeless mechanics. (This enables us to clearly separate out various components of the TTH which are easily conflated.) In §3-4, I explore two

[1] Although this problem already arises as an interpretive puzzle in classical theories like general relativity, the clash between treating diffeomorphisms as gauge symmetries and standard quantization procedures transforms the puzzle into a deep conceptual challenge for quantum theories of gravity. There is an extensive literature on the problem of time. For surveys, see Belot (2005), Thébault (forthcoming), and the references therein.

[2] See Swanson (2014, Ch. 2) for a philosophically-oriented introduction to modular theory and Borchers (2000) for a more detailed mathematical survey focusing on physical applications.

technical challenges concerning the relationship between thermal time and proper time conjectured by the TTH and the possibility of implementing the TTH in classical theories. Finally, in §5, I examine a trio of deeper conceptual problems concerning the flow of time in non-equilibrium states and the extent to which the TTH is background-independent and gauge-invariant. The outlook is mixed. I argue that while there are potentially viable strategies for addressing the two technical challenges, the three conceptual problems present a tougher hurdle for the defender of the TTH.

2 The Thermal Time Hypothesis

We usually think of theories of mechanics as describing the evolution of states and observables through time. Rovelli (2011) advocates replacing this picture with a more general timeless one that conceives of mechanics as describing relative correlations between physical quantities divided into two classes, *partial* and *full* observables. Partial observables are quantities that physical measuring devices can be responsive to, but whose value cannot be predicted given a state alone. Proper time along a worldline is an example of a partial observable. A clock carried by an astronaut measures their proper time, but the question "What is the astronaut's proper time?," is ill-posed. In contrast, a full observable is a coincidence or correlation of partial observables whose value can be determined given a state. Proper time along a worldline between points where it intersects two other wordlines is an example of a full observable. The question, "What is the proper time our astronaut experiences between launching from Earth and landing on the moon?," is well-posed and can be determinately answered once a state is specified. In general, only measurements of full observables can be directly compared to the predictions made by a mechanical theory.

For Hamiltonian systems, Rovelli's proposed timeless mechanics takes the following form: $\mathcal{C}$ is a configuration space of partial observables. A *motion* of the system is given by an unparametrized curve in $\mathcal{C}$, representing a sequence of correlations between partial observables. The dynamically possible motions are determined by the equation $H = 0$, where $H : T^*\mathcal{C} \to \mathbb{R}$ is a suitably smooth function, the *timeless Hamiltonian*, and the cotangent bundle, $T^*\mathcal{C}$, represents the space of partial observables and their conjugate momenta.[3]

[3]This is a special case of Rovelli's more general presymplectic framework for timeless

If the system has a privileged time variable, the Hamiltonian can be naturally decomposed into

$$H = p_t + H_0(q^i, p_i, t) \ , \tag{1}$$

where t is the partial observable in $\mathcal{C}$ that corresponds to time, and p_t is its conjugate momentum, i.e., the energy. The value of t can then be used to naturally parametrize the motions of the system in a manner similar to standard Hamiltonian mechanics.

In generally covariant systems, however, there is no privileged time variable, and the Hamiltonian lacks a canonical decomposition of the form (1). Although these systems are fundamentally timeless, it is possible for a notion of time to emerge thermodynamically. According to the second law of thermodynamics, a closed system will eventually settle into a thermal equilibrium state. Such states possess a range of unique properties. They are invariant with respect to the flow of time, stable under perturbations, and passive (i.e., mechanical work cannot be extracted by cyclic processes). Viewed as part of a definition of equilibrium, this thermalization principle requires an antecedent notion of time. The TTH inverts this definition and uses the notion of an equilibrium state to select a partial observable in $\mathcal{C}$ and identify it as time.

Suppose we know the full microstate of a generally covariant Hamiltonian system. Since the fundamental dynamics are given by the timeless Hamiltonian, and no special time variable is singled out, the flow of time is absent in our description of the system at the fundamental level. Suppose instead that we have a coarse-grained description of the system. If we somehow knew that this coarse-grained state were an equilibrium state, we could go on to identify the 1-parameter group of state space automorphisms with respect to which it is invariant, stable, and passive. The TTH conjectures that this interpretive move is always available. Given an arbitrary coarse-grained, statistical state it is possible to find a privileged 1-parameter group of state space automorphisms with respect to which it is in equilibrium. In this sense, the flow of time is a local, emergent phenomenon arising from our ignorance of the system's full state. Rovelli (2011) comments:

> When we say that a certain variable is "the time," we are not making a statement concerning the fundamental mechanical structure

mechanics.

> of reality. Rather, we are making a statement about the statistical distribution we use to describe the macroscopic properties of the system that we describe macroscopically. (p. 8)

This is the theoretical motivation for the TTH. While no partial observable plays a privileged temporal role in the fundamental physics, any coarse-graining will give rise to an emergent equilibrium dynamics that naturally selects a time parameter.

In practice, three hurdles present themselves. The first is providing a coherent mathematical characterization of equilibrium states. The second is finding a method for extracting information about the associated time flow from a specification of the coarse-grained 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 relevant limits.

For generally covariant quantum theories, Connes and Rovelli (1994) propose a concrete strategy to overcome these hurdles. Inspired by algebraic formulations of quantum field theory in curved spacetime, they propose treating a generally covariant quantum theory as a non-commutative C^*-algebra of diffeomorphically-invariant observables, $\mathfrak{A}$, along with a collection of physically possible states, $\{\rho\}$.[4] The states are (positive, normalized) linear functionals, $\rho : \mathfrak{A} \to \mathbb{C}$, encoding the expectation values of the observables in $\mathfrak{A}$. Choosing a state allows us to naturally expand $\mathfrak{A}$ to form a von Neumann algebra and represent it concretely as an algebra of bounded linear operators acting on a Hilbert space. This extra step is both physically and mathematically important, needed to handle boundary conditions in infinite quantum systems and to develop the tools of Tomita-Takesaki modular theory, but our discussion here will not hinge on the details.

In this setting, the properties of equilibrium states are captured by the *Kubo-Martin-Schwinger (KMS) condition.* Letting α_t be a 1-parameter group of automorphisms of $\mathfrak{A}$ representing the dynamics in the Heisenberg picture and $\beta = 1/T$ the inverse temperature, the condition requires that,

$$\rho(\alpha_t(A)B) = \rho(B\alpha_{t+i\beta}(A)) \tag{2}$$

for all $A, B \in \mathfrak{A}$. The left-hand side, $\rho(\alpha_t(A)B)$, represents the correlation between an arbitrary time-evolved observable, A, and another arbitrary observable, B, in the state ρ. In general, $\rho(\alpha_t(A)B) \neq \rho(B\alpha_t(A))$, but for

[4] See Brunetti et al. (2003) for a development of this basic theoretical framework.

equilibrium states, the KMS condition says that these re-ordered correlations can be obtained from one another by substituting $t \mapsto t + i\beta$.[5] The physical significance of this fact is not immediately obvious, but it turns out that (2) guarantees that ρ is stable, passive, and invariant under α_t. It is also a consequence of the more familiar Gibbs postulate,

$$\rho = \frac{e^{-\beta H}}{Tr[e^{-\beta H}]} \ , \tag{3}$$

which can be used to characterize equilibrium states in finite quantum systems governed by a discrete Hamiltonian, H. In order for (3) to be well-defined, the trace in the denominator must be finite for all values of β, but in infinite quantum systems where H can have a continuous spectrum (or might not exist at all), this cannot be guaranteed. The KMS condition, however, remains valid, and thus provides a mathematically tractable generalization of the Gibbs postulate for infinite quantum systems.[6]

The KMS condition allows Connes and Rovelli to clear the first hurdle. In order to clear the second, they appeal to the technical machinery of *Tomita-Takesaki modular theory.* A state is said to be *faithful* if $\rho(A^*A) = 0$ entails that $A = 0$. Since every non-zero observable has non-zero expectation value, a faithful state retains information about the full algebra, $\mathfrak{A}$. In this case, the tools of modular theory can be applied.[7] Its main theorem guarantees the existence of two unique modular invariants that depend on ρ, an antiunitary operator, J, and a positive operator, Δ. (Here we will only be concerned with the latter.) The theorem goes on to say that the set of unitary operators, $\{\Delta^{is} | s \in \mathbb{R}\}$, forms a (strongly continuous) 1-parameter group,

$$\sigma_s(A) := \Delta^{is} A \Delta^{-is} \ , \tag{4}$$

[5] In order for this substitution to make sense, there must exist a complex function, $F_{A,B}(z)$, analytic on the strip $\{z \in \mathbb{C} | 0 < \text{Im} z < \beta\}$ and continuous on the boundary of the strip, such that $F_{A,B}(t) = \rho(\alpha_t(A)B)$ and $F_{A,B}(t + i\beta) = \rho(B\alpha_t(A))$ for all $t \in \mathbb{R}$.

[6] See Bratteli and Robinson (1981, ch. 5.3-4) for a thorough introduction to the physics of KMS states, including their connection to the Gibbs postulate and proofs of stability and passivity properties.

[7] The traditional mathematical setting for modular theory involves a von Neumann algebra acting on a Hilbert space in standard form with respect to a cyclic, separating vector. The faithfulness assumption guarantees that these conditions will be met when $\mathfrak{A}$ is enlarged to form a von Neumann algebra in the canonical GNS representation defined by ρ.

called the *modular automorphism group.* The defining state is invariant under the flow of the modular automorphism group, $\rho(\sigma_s(A)) = \rho(A)$. Furthermore, it can be shown that $\rho(\sigma_s(A)B) = \rho(B\sigma_{s+i}(A))$. Thus ρ satisfies the KMS condition relative to $\{\sigma_s\}$ for inverse temperature $\beta = 1$.

This is Connes and Rovelli's central observation: in a generally covariant quantum theory, every faithful state determines a canonical 1-parameter group of automorphisms according to which it is a KMS state. This singles out a partial observable, the thermal time, $t_\rho := s$, parametrizing the flow of the (unbounded) thermal hamiltonian, $H_\rho := -\ln \Delta$. We can go on to decompose the timeless Hamiltonian, $H = p_{t_\rho} + H_\rho$. Associated with any faithful state, there is a natural "flow of time" according to which the system is in equilibrium.

It should be emphasized at this stage that the modular machinery employed by Connes and Rovelli requires ρ to be a mixed state. (Non-trivial C^*-algebras algebras have no pure, faithful states.) Giving mixed states an ignorance interpretation serves to connect their procedure to the guiding idea that a coarse-grained state determines the flow of time. According to this interpretation, it is because ρ is missing some information about the universe that a 1-parameter group of automorphisms naturally emerges. A state of maximal information would be pure, and the same machinery could not be used to explain the emergence of dynamics.

There are two important caveats here. First, not every mixed state is faithful, so Connes and Rovelli's proposal does not vindicate the idea that *any* coarse-graining determines a flow of time.[8] Second, as Wallace (2012) argues, nothing in the quantum formalism forces us to give mixed states an ignorance interpretation. On this alternative reading, the full microstate of the system could very well be faithful, and the flow of time, while still arising from the unique statistical features of such states, would no longer be a product of our ignorance. Thus there is a gap between much of the

[8] Why should we assume that typical coarse-grainings will be? In relativistic quantum field theory, the Reeh-Schlieder theorem ensures that the restriction of any global state analytic for the energy to any region whose causal complement is non-empty will be faithful. But this theorem relies on an antecedent specification of the dynamics as well as the background spacetime structure. In an arbitrary timeless mechanical theory these resources are unavailable. In light of this, the defender of the TTH might appeal to the following argument: insofar as we have reason to believe that relativistic quantum field theory is a good approximation of our world at some scale (and that the assumptions of the Reeh-Schlieder theorem apply to such effective field theories), we have reason to believe that our local statistical description of reality at that scale will be a faithful state.

motivational rhetoric surrounding the TTH and the current technical results. Based on the latter, all we can say is that given a statistically suitable (i.e., faithful) state, an equilibrium dynamics naturally emerges, whether or not this is a product of informational ignorance.

Regardless, the third hurdle remains. It is crucial to ask how this thermal time flow corresponds to various notions of physical time. In particular, how is thermal time related to the proper time measured by a localized observer? Although they do not establish a general theorem linking thermal time to proper time, Connes and Rovelli do make progress on the third hurdle in one important special case. Consider an immortal, uniformly accelerating observer in Minkowski spacetime. Because of their acceleration, they are only causally connected to a subregion of spacetime known as the *Rindler wedge*. The *Unruh effect*, a well-known, if rather mysterious physical effect predicts that our Rindler observer will measure a non-zero temperature in the vacuum state,

$$T_U = \frac{\hbar a}{2\pi k_B c} \tag{5}$$

where a is the magnitude of the observer's acceleration, k_B is Boltzmann's constant, $\hbar$ is the reduced Planck constant, and c is the speed of light. An inertial observer, in contrast, is causally connected to the entirety of Minkowski spacetime and measures $T_U = 0$ in the vacuum state.

Connes and Rovelli note that another deep theorem connecting modular theory to spacetime physics, the *Bisognano-Wichmann theorem*, provides a direct link between thermal time and the proper time measured by a Rindler observer. Let $\mathfrak{A}(W)$ be the algebra of observables localized in the Rindler wedge, W. The vacuum state is faithful for $\mathfrak{A}(W)$, so there is a well-defined modular automorphism group, $\{\sigma_s\}$, associated with vacuum wedge algebra. The Bisognano-Wichmann theorem says that $\{\sigma_s\}$ coincides with the group of wedge-preserving Lorentz boosts. Since the latter implement proper time translations along the Rindler observer's worldline, we find that thermal time is directly proportional to their proper time,

$$\frac{s}{\tau} = -T_U \ , \tag{6}$$

leading Connes and Rovelli to propose that the Unruh temperature should be physically interpreted as the ratio between thermal and proper time.[9]

We can now summarize the main content of the TTH:

[9]Working in units where $K_B = \hbar = c = 1$ for simplicity, the Bisognano-Wichmann

Thermal Time Hypothesis (Rovelli-Connes). *In a generally covariant quantum theory, given any faithful state, 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 their proper time.*

The TTH has three broad pillars: (I) the motivating idea that the flow of time is selected at the level of statistical mechanics in a fundamentally timeless, generally covariant theory, (II) a quantum mechanical model for such a selection mechanism, identifying thermal time with the modular automorphism group naturally associated with any faithful state, and (III) 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. Over the next three sections, we will consider a series of technical and conceptual objections to the TTH.

3 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. Consequently most of the attention that the TTH has received in the mathematical physics literature has focused on the third pillar noted above. But the Rindler observer is highly idealized, they are immortal and uniformly accelerating. Moreover, the Bisognano-Wichmann theorem assumes that the background global state is the vacuum state. Loosening each of these assumptions leads to technical complications which collectively appear quite daunting.

Realistic observers are mortal and therefore have causal access to a different region of spacetime, the doublecone formed by the intersection of their future lightcone at birth and past lightcone at death. What relationship holds

theorem entails that $\Delta = e^{-2\pi L_1}$, where L_1 is the generator of Lorentz boosts in the x^1-direction. The flow of proper time for an observer with uniform acceleration a in the x^1-direction is therefore given by the adjoint action of $e^{iL_1 a\tau}$. The modular automorphism group is given by the adjoint action of Δ^{is}, and so it follows that $s = -\frac{a}{2\pi}\tau$. (The minus sign is not physically significant. It results from the conventional mathematical definition of the modular automorphism group and can be eliminated by redefining σ_s as the adjoint action of Δ^{-is}.)

between thermal time and proper time for a uniformly accelerating, mortal observer in the vacuum state? An immediate problem arises due to the fact that their proper time is bounded whereas thermal time is unbounded. We therefore cannot expect a simple, linear relationship like (6).

Since the Rindler wedge can be related to a doublecone by a conformal transformation, in conformal field theories the modular group for the Rindler observer can be used to explicitly compute the modular group for their mortal counterpart. Using this trick, Martinetti and Rovelli (2003) prove that the mortal analogue of (6) is given by,

$$\frac{ds(\tau)}{d\tau} = \frac{-\hbar L a^2}{2\pi K_B c^3 \left(\sqrt{1 + \frac{a^2 L^2}{c^4}} - \cosh \frac{a\tau}{c} \right)} \,, \tag{7}$$

where L is half the lifetime of the observer who is born at $\tau = -L/c$ and dies at $\tau = L/c$. For both ling-lived and highly accelerating mortal observers, (7) converges to the Unruh temperature. Moreover, (7) is approximately constant for most of their lifespan, allowing the temperature to be interpreted as the local ratio between thermal and proper time in accordance with the TTH. Towards their birth and death, however, the two quantities rapidly diverge, calling this interpretation into doubt and raising a host of phenomenological questions.[10]

Unless our world is conformally-invariant, Martinetti and Rovelli's result is of limited applicability. Existing results strongly suggest that doublecone modular groups in generic quantum field theories cannot usually be given a dynamical interpretation. Any local dynamics must map doublecone-shaped subregions onto doublecone-shaped subregions while preserving spacelike separation and timelike ordering. Trebels (1997) proves that if a doublecone modular group satisfies these minimal geometric requirements, it will be re-

[10]If the observer's phenomenology is directly sensitive to thermal time rather than proper time, a straightforward reading of the physics suggests that at birth and death they experience a moment of infinite duration. While somewhat spiritually reassuring, this is rather physically implausible. If the observer directly experiences the flow of proper time instead, the TTH must explain how this experience emerges from the background thermal dynamics. A second issue: given that the temperature measured by a mortal observer is dependent on L, can the observer determine the date of their death by carefully measuring this temperature? Martinetti (2007) proves that the Heisenberg uncertainty principle prevents this. A finite observer will not live long enough to determine the temperature with the required accuracy.

lated to the conformal case by a scaling transformation.[11] In other words, a relationship like (7) is the best that we can hope for. But nothing guarantees that the doublecone modular groups in a generic quantum field theory will have these geometric properties, and Saffary (2005) has argued that in any theory with massive particles, they will not.[12]

The situation only becomes thornier when non-uniform acceleration and non-vacuum states are taken into account. In the latter case, it can be shown that the action of the modular automorphism group for the Rindler wedge in a non-vacuum state will differ from the vacuum action by a non-trivial automorphism.[13] This suggests that we should not generally expect the wedge modular group to have a geometric interpretation either. Even when it does, it will not be related to the vacuum action by some simple rescaling. In the case of a non-uniformily accelerating observer, the flow of proper time will fluctuate with acceleration while the flow of thermal time remains constant. If the TTH is correct, the observer will measure a shifting temperature reflecting the changing ratio between thermal and proper time (an idea supported by limited existing results on the Unruh effect for non-uniformly accelerating observers, e.g., Jian-yang et al. 1995). The challenge for the TTH is to explain the phenomenological experience of the observer, who will presumably age along with their fluctuating proper time, not the constant thermal time flow.

In the face of these obstacles, the defender of the TTH has at least four options on the table.

They can hold out hope for a suitably general dynamical interpretation

[11]For a detailed summary of Trebels's thesis work, see Borchers (2000, §3.4).

[12]In all known cases where they doublecone modular groups act geometrically, the group generators are ordinary differential operators of order one. In the known cases where they do not, Saffary proves that the generators contain a pseudodifferential term of order less than one. He goes on to argue that such terms typically give rise to non-local action, ruining any hopes of a geometric interpretation. More recently, Brunetti and Moretti (2010) have shown that in theories with massive particles the doublecone modular generators contain a pseudodifferential term of order zero. While the geometric ramifications of this fact have yet to be fully explored, combined with Saffary's analysis, it presents a major roadblock for extending Connes and Rovelli's original proposal to a wider class of mortal observers.

[13]The Radon-Nikodym theorem ensures that the action of the modular automorphism group uniquely determines the generating state. If ρ, ψ are two faithful states on a von Neumann algebra, then the associated modular automorphism groups $\{\sigma_s^\rho\}$, $\{\sigma_s^\psi\}$ differ by a non-trivial inner automorphism, $\sigma_s^\rho(A) = U\sigma_s^\psi(A)U^*$.

of modular groups for wedges and doublecones in a wide class of physically significant states. There is some indication that states of compact energy (including states satisfying the physically motivated Döplicher-Haag-Roberts and Buchholz-Fredenhagen selection criteria) give rise to well-behaved modular structure on wedges (Borchers, 2000). It is not clear that this is sufficient to ensure that modular automorphisms act geometrically, however, and in light of the limitations imposed by Trebels's and Saffary's no-go results, this first strategy seems like a long shot.

Alternatively, they could reject the idea that the thermal time flow determines the temporal metric directly. Thermal time would only give rise to the order, topological, and group theoretic properties of physical time. Metrical properties would be determined by a completely different set of physical relations. Some support for this idea comes from the justification of the clock hypothesis in general relativity. Rather than stipulating the relationship between proper time, τ, and the length of a timelike curve $||\gamma||$, Fletcher (2013) shows that for any $\epsilon > 0$, there is an idealized lightclock moving along the curve which will measure $||\gamma||$ within ϵ. This justifies the clock hypothesis by linking the metrical properties of spacetime to the readings of tiny idealized light-clocks. If 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 their proper time and not the geometrical flow of thermal time.

In line with this idea, Rovelli makes a number of allusions to the concept of an entropy clock as discussed by Eddington (1935). Eddington maintains 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 bifurcated.

A third strategy would be to argue that the metrical properties of time emerge from modular dynamics in the short-distance/high-energy limit. If a quantum field theory has a well-defined ultraviolet completion, the renormalization group flow in this regime should approach a conformal fixed point. Buchholz and Verch (1995) prove that in this limit, the doublecone 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 directly 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.

A final option would be to go back to the drawing board. Rovelli and Connes briefly note that since the modular automorphism groups associated with each faithful state are connected by inner automorphisms, they all project down onto the same 1-parameter group of outer automorphisms of the algebra.[14] The TTH could be revised to claim that this canonical state-independent flow represents the non-metrical flow of physical time.

It is unclear how viable of a strategy this really is. One significant reason for doubt is the fact that the move only works to recover a flow of time in quantum systems described by certain kinds of algebras, specifically type III (and type II_∞) von Neumann algebras. Only in these cases will the relevant 1-parameter group of outer automorphisms be non-trivial. In systems described by type I (and type II_1) algebras there is simply no passage of time according to the revised TTH.[15]

Although various theorems in algebraic and constructive quantum field

[14]An automorphism is *inner* if it is implemented by the adjoint action of a unitary element of the algebra. An *outer* automorphism is an equivalence class of automorphisms that can be related to each other by inner automorphisms. In general, the modular automorphism group for a given algebra and faithful state will not be inner and hence determines a non-trivial 1-parameter flow in the space of outer automorphisms. The Radon-Nikodym theorem, however, ensures that all of the modular automorphism groups over a given von Neumann algebra are inner-equivalent, and thus determine the same group of outer automorphisms.

[15]Von Neumann algebras can be classified as either type I, II_1, II_∞, or III based on their lattice of projection operators. In type I and type II_1 algebras, all modular automorphism groups are inner, hence their image in the group of outer automorphisms is trivial.

theory indicate that the local algebras assigned to doublecones and wedges are generically type III, it is not obvious that this will save the revised TTH. The global algebras that appear in quantum field theory are almost always type I, and so the revised TTH lacks a coherent story about the emergence of time at the cosmological level. Moreover, it is far from evident that observers like us actually have experimental and phenomenological access to the type III character of the local algebras at all. At the end of the day, when we make numerical calculations and do experiments, we use effective field theories and non-relativistic quantum mechanics to describe the world around us. In both cases we rely on type I algebras. The moral at the heart of the TTH is supposed to be that time emerges from a coarse-grained description of a fundamentally timeless reality. A moments reflection on the kinds of coarse-grained descriptions that we actually give renders the idea that the flow of time originates in the type III character of wedge and doublecone algebras rather implausible. It would be amazing if our experience of time had such a delicate source.

Due to these difficulties, it appears that the second strategy outlined above offers the best path forward for the defender of the TTH. Temporal topology and ordering is determined by the state-dependent modular automorphism group, while the temporal metric has a different origin, yet to be explained. Although this requires either modifying or abandoning the third pillar of the TTH, it preserve the first two pillars, and appears to be more plausible than a fully geometric interpretation of the local modular automorphism groups.

4 Thermal Time in Classical Theories

Turning attention to the first two pillars, it is interesting to note that the motivating idea that coarse-graining determines a thermal dyanmics does not obviously require the underlying timeless theory to be quantum mechanical. The proposed modular selection mechanism does, however, appear to crucially rely on the noncommutativity of quantum observables. Classical observables are typically represented by smooth functions over a state space manifold, and these can naturally be equipped with the structure of a commutative C^*-algebra. But in this case, since all observables commute, every state over the algebra is tracial, i.e., $\rho(fg) = \rho(gf)$ for all observables f, g. As a consequence of the main Tomita-Takesaki theorem, it follows that every modular automorphism group acts as the identity, trivializing the flow of thermal time.

Without a coherent mathematical procedure for extracting the thermal dynamics from a chosen state, a classical version of the TTH remains out of grasp. Does one exist, or is the TTH a uniquely quantum mechanical solution to the problem of time?[16] Investigating this question will help us to better understand the scope and content of the TTH. In addition, it could play a significant role in explaining the emergence of time in the classical limit, $\hbar \to 0$. Of course, a full understanding of this limit requires grappling with a host of entangled philosophical issues, most notably the measurement problem. Our aim here is more modest, to asses a proposed classical selection mechanism briefly sketched by Connes and Rovelli in their original paper. As we will go on to see, the idea stands on firmer foundational footing than one might initially suspect, and can even be linked to a classical analogue of Tomita-Takesaki modular theory.

Arguing by analogy with standard quantization procedures, Connes and Rovelli suggest that in classical theories commutators should be replaced by Poisson brackets. With respect to the Poisson bracket, the classical observables form a non-commutative algebra. Given an arbitrary coarse-grained state, ρ, represented by a probability distribution over state space, and reasoning by analogy with the Gibbs postulate, they introduce the "thermal Hamiltonian,"

$$H_\rho := -\ln \rho \ . \tag{8}$$

If ρ is nowhere vanishing (an assumption analogous to faithfulness in the quantum case), (8) defines a corresponding Hamiltonian vector field parametrized by thermal time, s. With respect to this vector field, the evolution of an arbitrary observable, f, is given by

$$\frac{d}{ds} f = \{-\ln \rho, f\} \ , \tag{9}$$

where $\{\ ,\ \}$ denotes the Poisson bracket.

In fact, the parallels between the classical and quantum versions of the TTH run much deeper than this ansatz hints. It is commonly thought that

[16]It should be noted that there is persistent disagreement over whether the problem of time itself is essentially quantum mechanical. Rovelli (2011) maintains that some version of the problem arises in any generally covariant theory, quantum or classical. For a recent dissenting viewpoint, see Pitts (2018).

the essential formal difference between classical and quantum mechanics is whether or not observables commute, but this emphasis on commutativity significantly obscures the rich algebraic structures employed by each theory.

On the quantum side, Alfsen and Shultz (1998) emphasize that the usual operator product in a non-commutative C^*-algebra is really two different products in disguise. It has a natural decomposition,

$$AB = A \bullet B - i(A \star B) \ , \tag{10}$$

where $A \bullet B$ is a commutative, nonassociative *Jordan product*, and $A \star B$ is a noncommutative, associative *Lie product*. The former, defined $A \bullet B := 1/2(AB + BA)$, encodes all spectral information about the observables. The latter, defined using the commutator, $A \star B := i/2(AB - BA)$, encodes the generating relationship between observables and state space symmetries. This observation reveals that non-commutative C^*-algebras are special cases of a more general class of *Lie-Jordan algebras*. Commutative C^*-algebras are not like this. Since the commutator vanishes, there is no natural Lie product, and essentially all that is leftover is a Jordan algebra encoding spectra.

On the classical side, Noether's theorem indicates that we should expect a similar generating relationship between observables and symmetries to hold. Indeed, in classical theories where state space is assumed to be a symplectic manifold, or more generally a Poisson manifold, it turns out that the algebra of observables also has a natural Lie-Jordan structure. Pointwise multiplication of smooth functions defines a commutative, associative Jordan product, $f \bullet g := fg$, encoding spectral information. The Poisson bracket determines a noncommutative, associative Lie product, $f \star g := \{f, g\}$, encoding how classical observables generate Hamiltonian vector fields.

The moral is this. If we naively choose to model classical systems using commutative C^*-algebras, we lose an important kind of information about the link between symmetries and observables. It is precisely this kind of information that is needed to formulate the technical details of the TTH. Reflecting on the physics, a much better choice is an associative Lie-Jordan algebra, which can be more directly compared to the nonassociative Lie-Jordan algebras employed by quantum theory.

This is the perspective adopted by the deformation and geometric quantization programs, two of the most mathematically rigorous approaches to quantization currently on the table.[17] In this setting, Gallavotti and Pul-

[17]See Landsman (1998) for an introduction to Lie-Jordan algebras and their role in

virenti (1976) use the Poisson bracket to define a classical analogue of the KMS condition, and Basart et al. (1984) link it to the quantum KMS condition in the $\hbar \to 0$ limit. This suggests that many of the tools needed for a classical version of the TTH already exist. Perhaps most compellingly, these include initial strides towards a classical analogue of Tomita-Takesaki modular theory made by Weinstein (1997).

Relative to a volume form, μ, on classical state space (assumed here to be a Poisson manifold), Weinstein defines the corresponding *modular vector field*,

$$\phi_\mu : f \to \mathrm{div}_\mu X_f \ , \tag{11}$$

where X_f is the Hamiltonian vector field associated with a classical observable, f. Weinstein proposes ϕ_μ as the classical analogue of the modular automorphism group. Intuitively, it characterizes the extent to which Hamiltonian vector fields, X_f, are divergence free with respect to the volume form μ, vanishing if and only if all Hamiltonian vector fields are divergence free.

Connecting the dots, we can trace a direct link between Weinstein's classical modular theory and the TTH. In the special case that state space is a symplectic manifold, there is a natural volume form, the *Liouville form*, defined in terms of the symplectic structure. Letting μ be the Liouville form, a quick calculation reveals that any nowhere-vanishing state, ρ, defines a non-trivial modular vector field,

$$\phi_{\rho\mu} = X_{-\ln\rho} \ , \tag{12}$$

equivalent to the vector field generated by the Hamiltonian $-\ln\rho$.[18] We immediately recognize this as the thermal Hamiltonian (8) postulated by Connes and Rovelli. The defining state, ρ, is invariant with respect to the corresponding dynamics and satisfies the Gibbs postulate, $\rho = e^{-\beta H_\rho}$, for inverse temperature $\beta = 1$. Just as in the quantum case, the classical thermal Hamiltonian can therefore be identified with the generator of state-dependent modular symmetries.

deformation and geometric quantization.

[18]In general, if h is a positive nowhere-vanishing function on a Poisson manifold, $h\mu$ defines a new volume form, and there is a simple expression relating the two modular vector fields, $\phi_{h\mu} = \phi_\mu + X_{-\ln h}$. In the symplectic case, if μ is the Liouville form, then $\phi_\mu(f) = 0$ for all observables f, since all Hamiltonian vector fields are divergence free with respect to μ. Since states are positive, it follows from putting these two facts together that ρ defines a modular vector field, $\phi_{\rho\mu} = X_{-\ln\rho}$.

This deep structural parallel suggests that the TTH is not essentially quantum mechanical and that the mathematics of modular theory can provide a coherent mechanism for selecting a preferred thermal time variable in classical theories too.[19] In addition, beyond its potential application to the TTH, Weinstein's framework offers a valuable laboratory for exploring the physical significance of modular theory in domains where the interpretational complexities bedeviling quantum theories do not arise.

5 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). Even if the third pillar needs to be modified in light of the challenges discussed in §3, the idea that a faithful state determines the non-metrical flow of time has proven resilient, and there is a plausible modular selection mechanism at play in both classical and quantum theories. There are, however, several deeper conceptual problems looming in the background which pose a more serious challenge to the hypothesis. Three of the most pressing raise questions about the coherence of the motivating idea behind the TTH and its adequacy in providing a solution to the problem of time.

The first is the *non-equilibrium problem*. While the TTH provides a coherent mathematical mechanism for selecting a non-metrical time flow, it is not clear that we should always view this flow as physical time. According to the thermal dynamics, the defining state is always a KMS state, but if it is a non-equilibrium state with respect to our ordinary conception of time, thermal time and physical time do not align. Relative to thermal time, a cube of ice in a cup of hot coffee is in an invariant equilibrium state! This is the "incredulous stare" that often confronts the TTH. Only for states which are true equilibrium states will the thermal time be physical time.

It would be incorrect to infer that the TTH rules out any thermodynamical change. A system in a KMS state can still exhibit fluctuations away from equilibrium. The defender of the TTH could try to argue that local non-equilibrium behavior can be viewed as fluctuations in some thermal

[19] There is an important caveat here: we lack classical analogue of the Bisognano-Wichmann theorem, so the third pillar of the TTH may turn out to be essentially quantum mechanical after all. Of course, the analysis in §3 indicates that this third pillar is rather shaky and will likely need some modification.

background state. On this approach, the local flow of time in my office according to which the ice melts and the coffee cools is not defined by the thermal state of the ice/coffee system, but the thermal state of some larger enveloping system.

Hints in this direction can be found in Rovelli (1993). In this paper, Rovelli explores the notion of thermal time in a spatially homogenous, isotropic Robertson-Walker universe filled with blackbody radiation. Such a model is a plausible approximation of a universe much like our own. The radiation represents the *cosmic microwave background (CMB)*, highly redshifted light left over from a phase early in the universe's history during which photons first decoupled from the cosmic plasma. In Robertson-Walker models, there is a natural notion of cosmological time given by the proper time experienced by a privileged class of observers co-moving with the expansion of the universe, for whom the universe always appears isotropic. (This is the time parameter usually employed in discussions of standard big bang cosmology.) Rovelli shows that the thermal time induced by equilibrium states of the CMB will be related to this cosmological time by a constant rescaling.

Although this non-trivial result is exciting, there remains a large explanatory gap between the physics it suggests and the temporal phenomenology of human observers. The CMB is a nearly uniform field of microwaves with temperature 2.7 Kelvin. Without specialized equipment — radio telescopes, radiometers, spectrophotometers — human observers would not even know it was there. It is highly implausible that our faculties of perception are sensitive to the thermal features of the CMB. The proposed chain of explanation must be longer and more complex. Thermal time explains the emergence of cosmological time, and then cosmological time is shown to be a natural measure of time for a particular class of observers. Even if the first half of this story can be convincingly filled out (see the related background-dependence problem introduced below), the vastly different scales involved in the second half of the story should give us pause. It is not exactly true that the universe appears isotropic to us. In only appears isotropic at the very largest scales, when we look beyond the Earth, the solar system, the Milky Way, the local group, etc. This process requires significant inductive extrapolation from what we directly experience.[20] At cosmological scales humans, stars, and

[20]Even the CMB does not appear isotropic to (sufficiently aided) human observers. The relative motion of the Earth in the CMB rest frame introduces a significant anisotropy in the CMB spectrum that must be factored out to reveal the usual images of a nearly smooth radiation field that we are familiar with.

galaxies might plausibly be viewed as small fluctuations in a largely homogeneous, isotropic background. Consequently, equilibrium thermal dynamics might well-describe the universe at this scale. But at the scales humans occupy, the local universe is highly non-homogneous and non-isotropic. If our temporal phenomenology is grounded in what we directly experience, and this experience is decidedly non-equilibrium, its hard to see how the order of explanation could plausibly run from cosmological time to local time.

Even if this challenge can be overcome, there is an additional wrinkle. Probably the most popular explanation for the arrow of time among physicists and philosophers alike, the *past hypothesis*, requires that in one temporal direction the universe is in an incredibly low-entropy state. But if thermal time is identified with physical time, this kind of asymmetric boundary condition is ruled out. The universe is in a KMS state with respect to the thermal dynamics. It has high entropy in both of the temporal directions determined by the flow of thermal time. The TTH is sometimes linked to the past hypothesis and motivated by parity of reasoning — if the direction of time has a thermodynamic origin, maybe the underlying flow of time does too — but the past hypothesis and the TTH are in fact deeply at odds with one another. The TTH forces us to adopt a rather unappealing "Boltzmann brain" view of cosmology as large-scale fluctuations from equilibrium.[21]

If the defender of the TTH balks at this conclusion, they have limited options on the table. One is to temper the view by only allowing certain reference states to determine the flow of thermal time, but the challenge of specifying a class of equilibrium states without an antecedent time flow was what prompted the permissiveness of the TTH in the first place. Furthermore, if a system is not actually in one of these reference states, it is hard to envision how a counterfactual state of affairs could determine the actual flow of time. This dilemma might motivate the defender of the TTH to explore the state-independent, outer modular flow as a last-ditch option. Identifying physical time with this flow would render it possible in principle to reconcile the TTH and the past hypothesis, however the criticisms discussed at the end of §3 must be overcome. In particular, if the global algebra is type I, then

[21] It might be possible to reconcile the TTH and the past hypothesis by treating the latter as a boundary condition for the observable universe, which is in turn viewed as a subsystem of a larger universe in thermal equilibrium. This move effectively embraces the Boltzmann brain cosmology one level higher up. Perhaps such a view will look more appealing situated within the landscape of a fundamentally timeless theory of quantum gravity. The jury is still out.

the triviality of the outer modular flow presents a new puzzle for quantum cosmology to grapple with.

A second, closely-related worry to the non-equilibrium problem has been voiced by Earman (2011) and Ruetsche (2014). In the physical situations where we can justify viewing the modular automorphism group as a kind of dynamics, it seems this is only possible because we already have a 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 scenario described by the original Bisognano-Wichman theorem, we are focused on spacelike wedges in Minkowski spacetime. We immediately recognize the geometric significance of the modular automorphism group because its flow is everywhere timelike. The orbits of σ_s correspond to a clear class of observer worldlines and $ds/d\tau$ is constant along those worldine, yielding a simple scaling relation between s and τ. Similarly, in the CMB model discussed above, the geometric interpretation of thermal time is secured by relating it to cosmological time in a highly symmetric Robertson-Walker universe. In other cases, even when the modular operators act geometrically, it can be hard to recognize the modular automorphisms as dynamical. The scaling relation for doublecone modular groups in conformal field theories (7) shows that the relation between thermal and proper time can be highly non-trivial (and this is the best case scenario for doublecones).

In general, calculating the explicit action of local modular automorphism groups is a very hard problem. In all of the cases outlined above where we can perform the calculations, extracting a dynamical interpretation requires antecedent knowledge of the background spacetime structure. In generic models of general relativity with no global timelike killing fields and no global isometries, such an interpretation may no longer be possible. Moreover if the ultimate goal is to use the TTH in conjunction with an eventual theory of quantum gravity to explain the emergence of spacetime itself, we cannot even appeal to the local Lorentzian geometry of spacetime to aid us. The problem is exacerbated if the TTH is modified in response to the non-equilibrium problem by restricting the set of states in which modular automorphisms define the flow of time. Unless the modular 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 suitably independent of background spatiotemporal structure. We will call this second problem, the *background-dependence problem.*[22]

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. The difficulty is that these partial observables are not diffeomorphism invariant. When an object changes position, we measure two gauge-invariant quantities, the position-of-the-object-at-s_1 and position-of-the-object-at-s_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 timeless Hamiltonian.

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 only diffeomorphism-invariant quantities will represent objective features of our world. The partial observables are just superfluous descriptive fluff. 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 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 there are no objective facts about absolute velocity in classical spacetime. There are however objective facts about relative velocities. By selecting a preferred reference frame (i.e., fixing a gauge), we can introduce absolute velocities into our theoretical description of the world and use them to compute gauge-invariant relative

[22]If the TTH is revised so that thermal time only generates the non-metrical properties of physical time, as suggested by the analysis in §3, the severity of the background-dependence problem is reduced, but only somewhat. Complicated scaling relations between thermal time and proper time cease to be an immediate issue, but the conditions under which local modular groups capture just the order, topological, and group theoretic properties of physical time are even less well understood. Plausibly, these conditions will depend on at least the conformal geometry of spacetime. On top of this, the defender of the TTH must also supply an entirely new explanation for the emergence of the temporal metric.

velocities. But we cannot use correlations between absolute velocities to *explain* facts about relative velocity. Objectively speaking, there are no such correlational facts to appeal to. Instead, we must restrict ourselves to the gauge-invariant structure of classical spacetime.

This structure can be captured by modeling classical spacetime as *Galilean spacetime*. Intuitively, Galilean spacetime consists of a time-ordered stack of 3-dimensional spatial slices. On each spatial slice, there is a metric determining facts about relative spatial distance between objects at that time. Across slices there is a temporal metric determining temporal distances and an affine structure characterizing deviations from inertial trajectories. Crucially, there is no spatial metric across slices, and as a result, there are no objective facts about absolute velocity. Despite this, there are still objective facts about relative velocity. Since the relative spatial distance between two objects at a given time is gauge-invariant, their relative velocity can be defined as the rate of change of this relative distance quantity. The worldlines of objects on in relative motion correspond to non-parallel 4-dimensional curves in Galilean spacetime.

Analogously, in a generally covariant setting we can freely 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 stage we 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.

A radical option is to reject the standard story about gauge symmetries. Rovelli (2014) 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. (p. 91)

On this picture, gauge-invariant quantities are intrinsically relational. Certain gauge-dependent quantities supply the relata, carrying modal information about the possible ways that free systems can interact with each other. For example, in classical electromagnetism, the interaction term in the Lagrangian, $-j^\mu A_\mu$, depends explicitly on the gauge-dependent vector potential. The Lagrangian itself is gauge-invariant, but to interpret it as describing the coupling between a charge density characterized by j^μ and the electromagnetic field, we must recognize A_μ as a genuine feature of the field, the "handle" to which charge couples (albeit in a manner that ultimately does not depend on the choice of a gauge-dependent coordinate system). The details of Rovelli's new proposal still need to be hammered out.[23] It should be emphasized that it marks a significant break from the received view on gauge.

Can a revised form of the TTH provide us with the explanatory tools to understand the flow of thermal time without reference to gauge-dependent partial observables, or does the framework of timeless mechanics require us to revise our conception of how explanation, ontology, and gauge symmetries are related? Whether or not the TTH can save time may ultimately rest on the solutions to these new reincarnations of vexingly familiar philosophical problems.

Acknowledgements

I would like to thank Carlo Rovelli, Jeremy Butterfield, David Baker, Gordon Belot, Hans Halvorson, Miklós Rédei, Laura Ruetsche, two anonymous referees, and audiences at Pittsburgh and the 2018 PSA meeting for helpful comments and criticism on previous versions of this paper.

References

Alfsen, E. and F. Shultz (1998). Orientation in operator algebras. *Proceedings of the National Academy of Sciences, USA 95*, 6596–6601.

Basart, H., M. Flato, A. Lichnerowicz, and D. Sternheimer (1984). Deformation theory applied to quantization and statistical mechanics. *Letters in Mathematical Physics 8*, 483–494.

Belot, G. (2005). The representation of time and change in mechanics. In J. Earman and J. Butterfield (Eds.), *Handbook of the Philosophy of Physics*. North-Holland.

[23]See Teh (2015) and Hetzroni (2020) for recent work in this direction.

Borchers, H. J. (2000). On revolutionizing quantum field theory with Tomita's modular theory. *Journal of Mathematical Physics 41*(6), 3604–3673.

Bratteli, O. and D. W. Robinson (1981). *Operator Algebras and Quantum Statistical Mechanics II.* Springer-Verlag.

Brunetti, R., K. Fredenhagen, and R. Verch (2003). The generally covariant locality principle – a new paradigm for local quantum field theory. *Communications in Mathematical Physics 237*, 31–68.

Brunetti, R. and V. Moretti (2010). Modular dynamics in diamonds. https://arxiv.org/abs/1009.4990.

Buchholz, D. and R. Verch (1995). Scaling algebras and renormalization group in algebraic quantum field theory. *Reviews in Mathematical Physics 7*, 1195.

Connes, A. and C. Rovelli (1994). Von Neumann algebra automorphisms and time-thermodynamics relation in generally covariant quantum theories. *Classical and Quantum Gravity 11*(12), 2899.

Earman, J. (2011). The Unruh effect for philosophers. *Studies in History and Philosophy of Modern Physics 42*, 81–97.

Eddington, A. (1935). *The Nature of the Physical World*. MacMilan.

Fletcher, S. (2013). Light clocks and the clock hypothesis. *Foundations of Physics 43*, 1369–1383.

Gallavotti, G. and M. Pulvirenti (1976). Classical KMS condition and Tomita-Takesaki theory. *Communications in Mathematical Physics 46*, 1–9.

Hetzroni, G. (2020). Gauge and ghosts. *British Journal for the Philosophy of Science*, https://doi.org/10.1093/bjps/axz021.

Jian-yang, Z., B. Aidong, and Z. Zheng (1995). Rindler effect for a nonuniformly accelerating observer. *International Journal of Theoretical Physics 34*, 2049–2059.

Landsman, N. (1998). *Mathematical Topics Between Classical and Quantum Mechanics*. New York: Springer-Verlag.

Martinetti, P. (2007). A brief remark on Unruh effect and causality. *Journal of Physics: Conference Series 68*, 012027.

Martinetti, P. and C. Rovelli (2003). Diamond's temperature: Unruh effect for bounded trajectories and thermal time hypothesis. *Classical and Quantum Gravity 20*(22), 4919.

Pitts, J. B. (2018). Equivalent theories and changing Hamiltonian observables in general relativity. *Foundations of Physics 48*, 579–590.

Rovelli, C. (1993). The statistical state of the universe. *Classical and Quantum Gravity 10*, 1567.

Rovelli, C. (2011). Forget time: Essay written for the FQXi contest on the nature of time. *Foundations of Physics 41*, 1475.

Rovelli, C. (2014). Why gauge? *Foundations of Physics 44*(1), 91–104.

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*. Ph. D. thesis, Hamburg.

Swanson, N. (2014). *Modular Theory and Spacetime Structure in QFT*. Ph. D. thesis, Princeton University.

Teh, N. J. (2015). A note on Rovelli's 'why gauge?'. *European Journal of the Philosophy of Science 5*, 339–48.

Thébault, K. (forthcoming). The problem of time. In E. Knox and A. Wilson (Eds.), *Companion to the Philosophy of Physics*. Routledge.

Trebels, S. (1997). *Über die Geometrische Wirkung Modularer Automoprhismen*. Ph. D. thesis, Göttingen.

Wallace, D. (2012). *The Emergent Multiverse*. Oxford University Press.

Weinstein, A. (1997). The modular automorphism group of a Poisson manifold. *Journal of Geometry and Physics 23*, 379–394.