Can Quantum Thermodynamics Save Time?

Noel Swanson*

January 30, 2019

Abstract

The *thermal time hypothesis (TTH)* is a proposed solution to the problem of time: a coarse-grained, statistical state determines a thermal dynamics according to which it is in equilibrium, and this dynamics is identified as the flow of physical 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 conjectured by the TTH and the implementation of the TTH in the classical limit. 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. 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.

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

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

translations correspond to a continuous 1-parameter subgroup of spacetime symmetries, and the dynamics are implemented either as a parametrized flow on state space (Schödinger 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 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. Given a privileged background time flow, one can characterize equilibrium states using various stability and passivity requirements. Conversely, given an equilibrium state, it is possible to derive the background dynamics. Rovelli (2011) exploits this converse connection, arguing that in a generally covariant theory, *any* coarse-grained, statistical state defines a notion of time according to which it is an equilibrium state. 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.[2] This paper represents a modest initial attempt to sally forth into rich philosophical territory. Its goal is to voice a number of impor-

[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. See Thébault (forthcoming) for a recent survey, Earman and Belot (2001), Earman (2002), Maudlin (2002), and Belot (2005) for influential philosophical discussions, and the second set of papers in Ashtekar and Stachel (1991) for a range of different perspectives from physicists.

[2]Earman (2002), Earman (2011), Arageorgis (2012), and Ruetsche (2014) are notable exceptions. Physicists have been more willing to dive in. For example, Paetz (2010) gives an excellent critical discussion of the many technical challenges faced by the TTH, some of which I explore in §3.

tant technical and conceptual problems faced by the TTH and to propose some strategies that the view has at its disposal to respond.

In §2, I provide a succinct 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 technical challenges concerning the relationship between thermal time and proper time conjectured by the TTH and the implementation of the TTH in the classical limit. 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 the state alone. A full observable is understood as a coincidence or correlation of partial observables whose value can be predicted given the state. Proper time along a worldline is an example of a partial observable. Proper time along a wordline between points where it intersects some other worldlines is an example of a full observable. Only measurements of full observables can be directly compared to the predictions made by the mechanical theory.

In Rovelli's timeless framework, a 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 state space 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 some appropriate vector space. For systems that can be modeled using

Hamiltonian mechanics, Γ and f are completely determined by a surface, Σ, in the cotangent bundle $T^*\mathcal{C}$ (the space of partial observables and their conjugate momenta p_a). This surface is defined by the vanishing of some Hamiltonian function $H : T^*\mathcal{C} \to \mathbb{R}$.

If the system has a preferred external time variable, the Hamiltonian can be decomposed as

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

where t is the partial observables in $\mathcal{C}$ that corresponds to time. A hallmark of generally covariant mechanical systems is that they lack such a canonical decomposition. Although these systems are fundamentally timeless, it is possible for a notion of time to emerge thermodynamically. By the second law of thermodynamics, a closed system 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. The TTH inverts 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 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 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. Minimally, such a theory can be thought of as a non-commutative C^*-algebra of diffeomorphism-invariant observables, $\mathfrak{A}$, along with a set of physically possible states, $\{\rho\}$.[3] Via the *Gelfand-Nemark-Segal (GNS) construction*, each state determines a concrete Hilbert space representation, $(\pi_\rho(\mathfrak{A}), \mathcal{H}_\rho)$, and a corresponding von Neumann algebra, $\pi_\rho(\mathfrak{A})''$, defined as the double commutant of $\pi_\rho(\mathfrak{A})$.

Next, Connes and Rovelli appeal to the well-known *Kubo-Martin-Schwinger (KMS) condition* to characterize equilibrium states. A state, ρ, on a von Neumann algebra, $\mathfrak{R}$, satisfies the KMS condition for inverse temperature $0 < \beta < \infty$ with respect to a 1-parameter group of automorphisms $\{\alpha_t\}$, if for any $A, B \in \mathfrak{R}$ 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,

[3]See Brunetti et al. (2003) for a development of this basic idea.

such that

$$F_{A,B}(t) = \rho(\alpha_t(A)B)$$
$$F_{A,B}(t + i\beta) = \rho(B\alpha_t(A)) \quad (2)$$

for all $t \in \mathbb{R}$. The KMS condition generalizes the idea of an equilibrium state to quantum systems with infinitely many degrees of freedom. KMS states are stable, passive, and invariant under the dynamics, $\{\alpha_t\}$. Moreover in the finite limit, the KMS condition reduces to the standard Gibbs postulate.[4]

Although the KMS condition is framed relative to a chosen background dynamics, according to the main theorem of *Tomita-Takesaki modular theory*, every faithful state determines a canonical 1-parameter group of automorphisms according to which it is a KMS state.[5] Connes and Rovelli go on to identify the flow of time with the flow of this state-dependent *modular automorphism group*.

In any GNS representation $(\pi_\rho(\mathfrak{A}), \mathcal{H}_\rho)$, the defining state, ρ, is represented by a cyclic vector, $\hat{\rho} \in \mathcal{H}_\rho$. This means that the norm closure of the subspace generated by $\pi_\rho(\mathfrak{A})\hat{\rho}$ is the entire GNS Hilbert space, $\mathcal{H}_\rho$. (Note that ρ is represented by a vector whether or not it is a pure or a mixed state.) If ρ is a *faithful* state, $\rho(A^*A) = 0$ entails that $A = 0$. In this special case the vector $\hat{\rho}$ is also separating, meaning that $\pi_\rho(A)\hat{\rho} = 0$ entails that $\pi_\rho(A) = 0$. In this setting, we can apply the tools of Tomita-Takesaki modular theory. Its main theorem asserts the existence of two unique modular invariants affiliated with $\pi_\rho(\mathfrak{A})''$, an antiunitary operator, J, and a positive operator, Δ. Here we will only be concerned with the latter. The 1-parameter family, $\{\Delta^{is} | s \in \mathbb{R}\}$, forms a strongly continuous unitary group,

$$\sigma_s(A) := \Delta^{is} A \Delta^{-is} \quad (3)$$

for all $A \in \pi(\mathfrak{A})''$, $s \in \mathbb{R}$. The defining state is invariant under the flow of the modular automorphism group, $\rho(\sigma_s(A)) = \rho(A)$. Furthermore, $\rho(\sigma_s(A)B) = \rho(B\sigma_{s-i}(A))$. Thus ρ satisfies the KMS condition relative to $\{\sigma_s\}$ for inverse

[4]See Ruetsche (2011, ch. 7.3) for a compact introduction to the physics of KMS states and Bratteli and Robinson (1981, ch. 5.3-4) for a detailed mathematical treatment.

[5]See Swanson (2014, ch. 2) for a brief philosophically-oriented introduction to modular theory, Borchers (2000) for a more detailed mathematical survey focusing on physical applications, and Kadison and Ringrose (1997, ch. 7) and Takesaki (2000) for comprehensive mathematical presentations.

temperature $\beta = 1$.[6]

For any faithful state, this procedure identifies a partial observable, the thermal time, $t_\rho := s$, parametrizing the flow of the (unbounded) thermal hamiltonian, $H_\rho := -\ln \Delta$, which has $\hat{\rho}$ as an eigenvector with eigenvalue zero. We can then go on to decompose the timeless Hamiltonian, $H = p_{t_\rho} + H_\rho$. Associated with any such faithful state, there is a natural "flow of time" according to which the system is in equilibrium.

It should be emphasized 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, statistical state determines the flow of time. On this reading, the flow of time is an emergent, local 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 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)

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* statistical state determines a flow of time.[7] Second, as Wallace (2012) argues, nothing in the quantum formalism forces us to give mixed states an ignorance

[6]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_s if and only if it is an arbitrary β KMS state with respect to $\sigma_{s'}$, where $s' = s/\beta$.

[7]Why should we assume that our macroscopic description of reality will be faithful? 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 (see Horuzhy 1990, thm. 1.3.1). 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. A generalization of Connes and Rovelli's procedure that eschews the assumption of faithfullness would be a welcome development. In its absence, the defender of the TTH could 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. While

interpretation. On this alternative reading, the full state of the system could very well be mixed, and the flow of time, while still arising from the unique statistical features of mixed states, is no longer a product of our ignorance.

Regardless of its origin, it is important 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. For a uniformly accelerating, immortal observer in Minkowski spacetime, the region causally connected to her worldline is the *Rindler wedge*. In standard coordinates we can explicitly write the observer's trajectory as

$$\begin{aligned} x^0(\tau) &= a^{-1}\sinh(\tau) \\ x^1(\tau) &= a^{-1}\cosh(\tau) \\ x^2(\tau) &= x^3(\tau) = 0 \end{aligned} \tag{4}$$

where τ is the observer's proper time. The wedge region is defined by the condition $x^1 > |x^0|$. The *Bisognano-Wichmann theorem* then tells us that in the vacuum state, the modular automorphism group for the wedge implements wedge-preserving Lorentz boosts — Δ^{is} is given by the boost $U(s) = e^{2\pi i s K_1}$ (where K_1 is the generator of an x^1-boost). Since the Lorentz boost $\lambda(a\tau)$ implements a proper time translation along the orbit of an observer with acceleration a, $U(\tau) = e^{ai\tau K_1}$ can be viewed as generating evolution in proper time. Comparing these two operators, we find that proper time is directly proportional to thermal time,

$$s = \frac{2\pi}{a}\tau\ . \tag{5}$$

The Unruh temperature measured by the observer is $T = a/2\pi k_b$ (where k_b is Boltzmann's constant), this leads Connes and Rovelli to propose that the Unruh temperature can be interpreted as the ratio between thermal and proper time. Not only does this relationship hold along the orbits of constant acceleration, but as Paetz (2010) emphasizes, if an observer constructs global

this justification might suffice for some purposes, it would put significant restrictions on the scope of the TTH.

time coordinates for the wedge via the process of Einstein synchronization, this global time continues to coincide with the rescaled thermal time flow.[8]

We can summarize the main content of the TTH as follows:

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 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 state-dependent modular flow on the algebra of observables, 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, I 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. But the theorem only applies to wedge algebras in the

[8] 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 of constant η 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 $a/2\pi$.

vacuum state. 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.

We can attempt to generalize these results by considering (a) mortal observers, (b) non-uniformly accelerating observers, and (c) non-vacuum states. In the case of an immortal, 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 (a)-(c), if the modular automorphism group acts geometrically at all, there are two competing flows which are not related by a simple linear 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.

Starting with (a), a uniformly accelerating mortal observer has causal access to a different region of Minkowski spacetime, the doublecone formed by the intersection of her future lightcone at birth and her past lightcone at death. The relationship between thermal time and proper time in this case must be more complicated due to the fact that the the proper time experienced by a finite observer is bounded while the modular time is unbounded. Because wedges and doublecones can be related by a conformal transformation, in conformally-invariant theories, geometric results from wedge algebras can be transferred onto the doubelcone algebras. In this case, Martinetti and Rovelli (2003) prove that

$$s = \frac{2\pi}{La^2}(\sqrt{1 + a^2L^2} - \cosh a\tau) \ , \tag{6}$$

where L is half the lifetime of the observer. For most of the observer's lifespan, $s(\tau)$ is approximately constant, allowing the Unruh temperature to be interpreted as the local ration between thermal and proper time for such observers.[9] 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.

Unless our world is conformally-invariant, however, this result is of limited applicability. At present, there are no general results ensuring that the

[9]While the Unruh temperature for a non-accelerating immortal observer is zero, for mortal non-accelerating observer it 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.

modular operators of doublecones in a generic model of QFT have a geometric interpretation in the vacuum state. The most systematic investigation of this question has been done by Trebels (1997).[10] Let O be an open, connected, causally complete region of Minkowski spacetime, and let $\mathfrak{R}(O)$ be the local von Neumann algebra. A unitarily implemented automorphism, $U\mathfrak{R}(O)U^* = \mathfrak{R}(O)$, is *geometric, causal, and order preserving*, if there exists a 1-1 map, $g : O \to O$, such that

(i) for every doublecone $D \subset O$, $U\mathfrak{R}(D)U^* = \mathfrak{R}(g(D))$.

(ii) if $x, y \in O$ and $x - y$ is spacelike, then $g(x) - g(y)$ and $g^{-1}(x) - g^{-1}(y)$ are spacelike,

(iii) if $x, y \in O$ and $x - y$ is in the forward lightcone, then $g(x) - g(y)$ and $g^{-1}(x) - g^{-1}(y)$ are in the forward lightcone.[11]

If the modular automorphism group has a local dynamical interpretation, σ_s is a geometric, causal, order preserving automorphism, for every $s \in \mathbb{R}$. This will give rise to a corresponding group of mappings $g(s)$ satisfying conditions (i)-(iii) above. Using this definition, Trebels goes on to prove that for doublecone algebras in the vacuum state, if the associated modular operators act geometric, causal, and order preserving, they must act as scaled versions of the modular operators in a conformal theory. So equation (6) represents the most general possible geometric relationship possible between s and τ.[12]

Of course all of this hinges on σ_s having a local dynamical interpretation, which is not guaranteed. Saffary (2005) argues that in massive theories, we should not generally expect the doublecone modular automorphism group to act geometric, causal, and order preserving. In the massless case, the theory is conformally-invariant and the modular generators are ordinary differential operators, δ_0, of order one. In the massive case, it has long been conjectured

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

[11] Condition (i) capture the relevant sense in which the action is geometric, while (ii) and (iii), respectively, capture the senses in which the action is causal and order preserving.

[12] Using this result one can argue quite broadly 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 on Minkowski spacetime does not restrict to a local Killing field on the doublecone. In contrast, if the modular automorphism group acts geometric, causal, and order preserving on the doublecone, then it generates a local conformal Killing field (Paetz, 2010).

that the modular generators are pseudo-differential operators $\delta_m = \delta_0 + \delta_r$, where the leading term is given by the massless generator, and δ_r is a pseudo-differential operator of order less than one. The second term is thought to give rise to non-local action without geometric interpretation, but the formal results backing this argument are only partial at this stage.[13]

The case of non-uniformly accelerating observers (b) generates a second set of worries. Assuming that the observer is immortal, she is once again confined to the Rindler wedge. In the vacuum state, 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.[14] 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 is 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.

A third challenge is presented by the case of non-vacuum states (c). The

[13]In all known cases where the modular automorphisms act geometric, causal, and order preserving, the modular generators are ordinary differential operators of order one. Saffary proves that in the two known cases where they do not have a geometric interpretation (Yngvason, 1994; Borchers and Yngvason, 1999), the generators are of the form $\delta_m = \delta_0 + \delta_r$. More recently, Brunetti and Moretti (2010) have confirmed that δ_r is in fact a pseudo-differential operator of order zero. The geometric ramifications of this result have yet to be fully explored.

[14]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

$$T = \frac{\pm a(t)/(1 - 2\dot{x}_H)}{2\pi k_b} \tag{7}$$

for $a(t) > 0$, $a(t) < 0$ respectively. Here $a(t)$ is the acceleration, t the coordinate 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.

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 $\mathfrak{R}$, then the associated modular automorphism groups σ_s^ρ, σ_s^ψ differ by a non-trivial inner automorphism, $\sigma_s^\rho(A) = U\sigma_s^\psi(A)U^*$, for all $A \in \mathfrak{R}$, $s \in \mathbb{R}$. Thus for an immortal, 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 the vacuum modular automorphism group by rescaling.

None of these are knockdown objections since so little is known about the geometric action of modular operators apart from the Bisognano-Wichmann theorem and its generalization for conformal theories. But our current ignorance also presents a major challenge. (The situation is even less clear in general curved spacetime settings.) The defender of the TTH has at least four options on the table.

She can hold out hope for a suitably general dynamical interpretation of modular automorphisms in a wide class of physically significant states. There is some indication that states of compact energy (including states satisfying the Döplicher-Haag-Roberts and Buchholz-Fredenhagen selection criteria) give rise to well-behaved modular structure on wedges. In this case the wedge modular automorphisms can be related to those in the vacuum state by the Radon-Nikodym derivative (Borchers, 2000). (The analogous problem for doublecones is still open.) 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 overall strategy seems like a long shot.

Alternatively, she 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.[15] This justifies the clock hypothesis

[15] See also Maudlin (2012, Ch. 5).

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 her 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.

Perhaps motivated by the justification of the clock hypothesis, the defender of the TTH could attempt to argue that the metrical properties of time emerge from modular dynamics in the short distance limit of the theory. If the theory has a well-defined ultraviolet limit, the renormalization group flow 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 of a von Neumann algebra are connected by inner automorphisms, they all project down onto the same 1-parameter group of outer automorphisms of the algebra.[16] The TTH could be revised to claim that this canonical state-independent flow represents the non-metrical flow of physical time. It is not known under what circumstances the outer flow acts in a suitably geometric fashion to be interpretable as a dynamics, so it remains to be seen whether or not this is a viable option.

There is a significant reason for doubt — the move only works to recover a flow of time in systems described by type III (and type II_∞) von Neumann algebras. 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. For systems described by such algebras, there is simply no passage of time according to the revised TTH.[17]

Although various theorems in algebraic and constructive quantum field theory indicate that the local algebras of doublecones and wedges are type III it is not obvious that this will save the revised TTH. For one thing, these results all critically rely on the analytic properties of continuous spacetime translations, and in the case of doublecones, on the ultraviolet scaling properties of the quantum fields.[18] It is unknown if the type III character of local algebras is a physical effect that survives when quantum field theory is viewed as a low-energy approximation of quantum gravity. (Effective field theories are typically formulated using type I algebras.) To compound this problem, even if the type III property is a physical effect, it is highly questionable that macroscopic observers like us will be sensitive to it. The same theorems establishing the type III character of local algebras in quantum

[16]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.

[17]Every von Neumann algebra contains a complete lattice of projection operators whose structure can be used to classify the algebra as type I, II, or III. Non-relativistic quantum mechanics uses type I algebras almost exclusively, while type III algebras are generic in existing models of relativistic quantum field theory. See Ruetsche (2011, ch. 7) and Kadison and Ringrose (1997, ch.6) for details.

[18]See Halvorson and Müger (2006, §2.5) for a survey of these theorems.

field theory also indicate that they are hyperfinite, meaning that they can be approximated to arbitrary accuracy by funnels of finite type I algebras. Type I funnels play a central role in recent analyses of measurement in quantum field theory (Okamura and Ozawa, 2016). If our grip on local physics comes from such measurements, it does not appear possible for us to directly probe the type III character of our immediate environment. It would be amazing if our experience of the flow 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 (the interpretation of the Unruh temperature as the ratio between proper and thermal time), it preserve the first two pillars, and appears to be more plausible than a fully geometric interpretation of the local modular automorphism groups.

4 The Classical Limit

The classical limit presents a different kind of challenge. Prima facie, nothing about the idea that a statistical state selects a preferred thermal time requires that the theory be quantum mechanical. The proposed mechanism for selecting a partial observable 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, $\rho(AB) = \rho(BA)$ for all $A, B \in \mathfrak{R}$. Consequently, every modular automorphism group acts as the identity, trivializing the thermal time flow. Does the TTH have a classical counterpart, or is the flow of time ultimately a quantum mechanical phenomena?

Arguing by analogy with standard quantization procedures, Connes and Rovelli suggest that in the classical limit commutators need to be replaced by Poisson brackets. We begin with an arbitrary statistical state, ρ, represented by a probability distribution over classical state space, Γ:

$$\int_{\Gamma} dx\ \rho(x) = 1\ , \tag{8}$$

where $x \in \Gamma$ is a timeless microstate. By analogy with the Gibbs postulate, we can introduce the "thermal Hamiltonian,"

$$H_\rho := -\ln \rho \ , \tag{9}$$

which, if ρ is nowhere vanishing, defines a corresponding Hamiltonian vector field.[19] With respect to this vector field, the evolution of an arbitrary classical observable, $f \in C^\infty(\Gamma)$, is given by

$$\frac{d}{ds} f = \{-\ln \rho, f\} \ , \tag{10}$$

and $\rho = e^{-H_\rho}$. 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, and Basart et al. (1984) use tools from the deformation quantization program to link it to the quantum KMS condition in the $\hbar \to 0$ limit. Is this connection strong enough to support a version of the TTH in ordinary general relativity, or does it only serve to aid us in understanding how the thermal time variable behaves in the transition from an underlying quantum theory to emergent classical physics?

A significant 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.[20]

Setting aside these broader interpretive challenges for now, an important first step lies in obtaining a better understanding the classical selection procedure outlined above. As it turns out, the commutator-to-Poisson-bracket ansatz is on firmer foundational footing than one might initially suspect. As emphasized by Alfsen and Shultz (1998), non-abelian C^*-algebras have a natural *Lie-Jordan structure*:

$$AB = A \bullet B - i(A \star B) \ , \tag{11}$$

[19]The requirement that ρ be nowhere vanishing is analogous to the assumption of faithfulness in the quantum case.

[20]See Rovelli and Smerlak (2011).

The non-associative Jordan product, $\bullet$, encodes information about the spectra of observables, while the associative Lie product, $\star$, encodes the generating relation between observables and symmetries. The significance of the commutator, is that it defines the canonical Lie product, $A \star B := i/2[A, B]$. Classical mechanical theories formulated on either a symplectic or Poisson manifold have a natural Lie-Jordan structure as well. The standard product of functions defines an associative Jordan product, encoding spectral information, while the Poisson bracket determines the associative Lie product, describing how classical observables generate Hamiltonian vector fields on state space. Together, this structure is called a *Poisson algebra*. The primary difference between the classical and quantum cases is the associativity/non-associativity of the Jordan product.

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. In this setting, 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 the corresponding *modular vector field*, ϕ_μ, as the operator $\phi_\mu : f \to \mathrm{div}_\mu X_f$ where X_f is the Hamiltonian vector field associated with a classical observable, $f \in C^\infty(\Gamma)$.[21] The antisymmetry of the Poisson bracket entails that 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 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. If h is a positive function, $h\mu$ defines a new smooth density, and there is a simple expression relating ϕ_μ and $\phi_{h\mu}$:

$$\phi_{h\mu} = \phi_\mu + X_{-\ln h} \tag{12}$$

In the special case that Γ is a symplectic manifold and μ is the density associated with the canonical Liouville volume form, $\phi_\mu(f) = 0$ for all observables. This reflects the conservation of energy by Hamiltonian flows in symplectic dynamical systems. Any statistical state, ρ, can thus associated with a non-trivial modular vector field,

$$\phi_{\rho\mu} = X_{-\ln \rho} \ , \tag{13}$$

[21]Here the divergence of a vector field, ξ, is defined as $\mathrm{div}_\mu \xi := \mathcal{L}_\xi \mu / \mu$ where $\mathcal{L}_\xi$ is the Lie derivative with respect to ξ.

equivalent to the Hamiltonian vector field $X_{-\ln\rho}$ with Hamiltonian $-\ln\rho$. We immediately recognize this as the thermal Hamiltonian, H_ρ, 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. This is an intriguing parallel between the structure of classical and quantum mechanics and suggests that the mathematical machinery of modular theory can provide a coherent mechanism for selecting a preferred thermal time variable in classical theories too. If quantum mechanics turns out to be an essential component of the TTH, it seems that it will only be as part of the third pillar linking thermal time and proper time.

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 Unruh conjecture needs to be modified in light of the challenges discussed in §3, the idea that a statistical 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 viability of 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 can 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, the resulting 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 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 correspond to physical time. Of course without a 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.

It would be incorrect to infer that the TTH rules out any thermodynamics change. A system in a KMS state can still exhibit fluctuations away from thermal equilibrium. The defender of the TTH could try to argue that local non-equilibrium behavior can be viewed as 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 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 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. 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 observers like us. It would be truly remarkable to discover that our faculties of perception are sensitive to the thermal features of the CMB.

Even if local non-equilibrium physics can be successfully explained in terms of fluctuations, applying this strategy at the cosmological level is dubious. Probably the most popular explanation for the arrow of time among physicists and philosophers alike, the *past hypothesis* (Albert, 2003), 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. 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.[22]

If the defender of the TTH balks at this conclusion, she has limited options

[22]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.

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 addition, even in quantum field theories with type III local algebras, the global algebra is expected to be type I.[23] In this case there is no global flow of time, even if the state of the universe is a statistical one, giving rise to a new puzzle for quantum cosmology to grapple with.

A second, closely-related conceptual worry 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 Bisognano-Wichman scenario, 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 $d\tau/ds$ is constant along those worldine, yielding a simple scaling relation between s and τ. In other cases, even when the modular operators act geometrically, it can be hard to recognize σ_s as dynamical. The conformal example (6) shows 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 state 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 $x^\mu \mapsto e^{-2\pi s}x^\mu$ on vectors in V^+.[24] If we interpret the orbits of σ_s as the worldlines of inertial observers

[23]See, for example, Horuzhy 1990, prop. 1.3.46.
[24]Haag (1996, Thm. 4.2.2).

departing from the origin of the lightcone, the relationship between thermal time and proper time has the form $s(\tau) = -\ln \tau/2\pi$.

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 and no global isometries, such an interpretation may no longer be possible. 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.*

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. 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 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 velocities 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 structure of Galilean spacetime.[25]

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 explananda into the explanans. 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

[25]The spatial distance between two objects at a time is gauge invariant and we can define relative velocity as the rate of change of this relative spatial distance. Four-dimensionally, worldlines of objects in relative motion with one another in Galilean spacetime will be represented by non-parallel inertial trajectories.

in. But the explanatory role of the thermal time partial observable extends well beyond this minimal indexicality.

A more radical alternative is to reject the standard picture of gauge symmetry entirely. 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)

It should be emphasized that this marks a significant 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.

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 ontology, explanation, and gauge symmetries are related? Whether or not quantum thermodynamics can save time may ultimately rest on the solutions to these new reincarnations of vexingly familiar philosophical problems.

References

Albert, D. (2003). *Time and Chance*. Harvard University Press.

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

Arageorgis, A. (2012). The philosophy of physics meets the Tomita-Takesaki modular theory. Second Summer School on Operator Theory, National Technical University of Athens, July 23-28.

Ashtekar, A. and J. Stachel (Eds.) (1991). *Conceptual Problems of Quantum Gravity*, Volume 2 of *Einstein Studies*. Birkhäuser.

Basart, H., M. Flato, A. Lichnerowicz, and D. Sternheimer (1984). Deformation theory applied to quantizatio 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.

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 J. Yngvason (1999). Modular groups of quantum fields in thermal states. *Journal of Mathematical Physics 40*, 601–24.

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. (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 G. Belot (2001). Pre-socratic quantum gravity. In C. Callender and N. Huggett (Eds.), *Physics Meets Philosophy at the Planck Scale.*

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.

Haag, R. (1996). *Local Quantum Physics.* Berlin: Springer-Verlag.

Halvorson, H. and M. Müger (2006). Algebraic quantum field theory. In J. Butterfield and J. Earman (Eds.), *Philosophy of Physics*, pp. 731–922. Elsevier.

Horuzhy, S. S. (1990). *Introduction to Algebraic Quantum Field Theory.* Kluwer Academic Publishers.

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

Kadison, R. V. and J. R. Ringrose (1997). *Fundamentals of the theory of operator algebras. Vol. II.* American Mathematical Society.

Martinetti, P. (2004). A brief remark on Unruh effect and causality.

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.

Maudlin, T. (2002). Thoroughly muddled mctaggart. *Philosophers' Imprint 2*, http://www.philosophersimprint.org/002004/.

Maudlin, T. (2012). *Philosophy of Physics: Space and Time.* Princeton University Press.

Okamura, K. and M. Ozawa (2016). Measurement theory in local quantum physics. *Journal of Mathematical Physics 57*, 015209–29.

Paetz, T.-T. (2010). An analysis of the 'thermal-time concept' of Connes and Rovelli. Master's thesis, Georg-August-Universität Göttingen.

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

Rovelli, C. and M. Smerlak (2011). Thermal time and Tolman–Ehrenfest effect: 'temperature as the speed of time'. *Classical and Quantum Gravity 28*(7), 075007.

Ruetsche, L. (2011). *Interpreting Quantum Theories*. Oxford: Oxford UP.

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.

Takesaki, M. (2000). *Theory of Operator Algebras, Vol. II*. Springer-Verlag.

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.

Yngvason, J. (1994). A note on essential duality. *Letters in Mathematical Physics 31*, 127–141.