

To be presented at the *2018 PSA Meeting*:

Can Quantum Thermodynamics Save Time?

Noel Swanson*

Abstract

The *thermal time hypothesis (TTH)* is a proposed solution to the problem of time: every 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 implementation of the TTH in the classical limit and the relationship between thermal time and proper time. Two more conceptual problems focus on interpreting the flow of time in non-equilibrium states and the lack of gauge invariance.

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 statespace (Schödinger picture) or a parametrized group of automorphisms of the algebra of observables (Heisenberg picture). In generally

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

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 on different values at different times, also fails. This is the infamous *problem of time*.

Connes and Rovelli (1994) propose a radical solution to the problem: the flow of time (not just its direction) has a thermodynamic origin. Equilibrium states are usually defined with respect to a background time flow (e.g., dynamical stability and passivity constraints reference a group of time translations). Conversely, given an equilibrium state one can derive the dynamics according to which it is in equilibrium. Rovelli (2011) exploits this converse connection, arguing that in a generally covariant theory, *any* statistical state defines a notion of time according to which it is an equilibrium state. 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.¹ 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 problems faced by the TTH and to highlight some tools that the view has at its disposal to respond.

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

¹Earman (2002), Earman (2011), and Ruetsche (2014) are notable exceptions. Physicists have been more willing to dive in. Paetz (2010) gives an excellent critical discussion of the many technical challenges faced by the TTH.

given the state alone (e.g., proper time along a worldline). A full observable is understood as a coincidence or correlation of partial observables whose value can be predicted given the state (e.g., proper time along a worldline at the point where it intersects another worldline). Only measurements of full observables can be directly compared to the predictions made by the mechanical theory.

A timeless mechanical system is given by a triple (\mathcal{C}, Γ, f) . \mathcal{C} is the configuration space of partial observables, q^a . A *motion* of the system is given by an unparametrized curve in \mathcal{C} , representing a sequence of correlations between partial observables. The space of motions, Γ is the statespace of the system and is typically presymplectic. The evolution equation is given by $f = 0$, where f is a map $f : \Gamma \times \mathcal{C} \rightarrow V$, and V is a vector space. For systems that can be modeled using Hamiltonian mechanics, Γ and f are completely determined by a surface Σ in the 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 $H : T^*\mathcal{C} \rightarrow \mathbb{R}$.

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. Generally covariant mechanical systems lack such a canonical decomposition. Although these systems are fundamentally timeless, it is possible for a notion of time to emerge thermodynamically. A closed system left to thermalize will eventually settle into a time-independent equilibrium state. Viewed as part of a definition of equilibrium, this thermalization principle requires an antecedent notion of time. 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 as a non-commutative C^* -algebra of diffeomorphism-invariant

observables, \mathfrak{A} , along with a set of physically possible states, $\{\phi\}$.² Via the Gelfand-Nemark-Segal (GNS) construction, each state determines a concrete Hilbert space representation $(\pi_\phi(\mathfrak{A}), \mathcal{H}_\phi)$, and a corresponding von Neumann algebra $\pi_\phi(\mathfrak{A})''$, defined as the double commutant of $\pi_\phi(\mathfrak{A})$.

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

$$\begin{aligned} F_{A,B}(t) &= \rho(\alpha_t(A)B) \\ F_{A,B}(t + i\beta) &= \rho(B\alpha_t(A)) \end{aligned} \tag{2}$$

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.

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. Connes and Rovelli go on to identify the flow of time with the flow of this state-dependent *modular automorphism group*.

In the GNS representation $(\pi_\phi(\mathfrak{A}), \mathcal{H}_\phi)$, the defining state, ϕ , is represented by a cyclic vector $\Phi \in \mathcal{H}_\phi$. If ϕ is a *faithful* state (i.e., if $\phi(A^*A) = 0$ entails that $A = 0$) then the vector Φ is also separating. In this setting we can apply the tools of Tomita-Takesaki modular theory. The main theorem asserts the existence of two unique modular invariants, 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,

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

for all $A \in \pi(\mathfrak{A})''$, $s \in \mathbb{R}$. The defining state is invariant under the flow of the modular automorphism group, $\phi(\sigma_s(A)) = \phi(A)$. Furthermore, $\phi(\sigma_s(A)B) =$

²See Brunetti et al. (2003) for a formal development of this basic idea.

$\phi(B\sigma_{s-i}(A))$. Thus ϕ satisfies the KMS condition relative to $\{\sigma_s\}$ for inverse temperature $\beta = 1$.

For any faithful state, this procedure identifies a partial observable, the thermal time, $t_\phi := s$, parametrizing the flow of the (unbounded) thermal hamiltonian $H_\phi := -\ln \Delta$, which has Φ as an eigenvector with eigenvalue zero. We can then go on to decompose the timeless Hamiltonian $H = p_{t_\phi} + H_\phi$. Associated with any such state, there is a natural “flow of time” according to which the system is in equilibrium. But in what sense does this thermal time flow correspond to various notions of physical time? In particular, how is thermal time related to the proper time measured by a localized observer?

Although they do not establish a general theorem linking thermal time to proper time, Connes and Rovelli do make substantial progress on the third hurdle in one intriguing 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 $U(s) = e^{2\pi is K_1}$ (where K_1 is the representation of 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 if an observer constructs global time coordinates for the wedge via the process of Einstein synchronization, this global time continues to coincide with the rescaled thermal time flow.

We can now summarize the main content of the TTH:

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.*

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

The Bisognano Wichmann theorem only applies to immortal, uniformly accelerating observers in the vacuum state of a quantum field theory in flat spacetime. How can we characterize the relationship between thermal and proper time for a broader, more physically realistic class of observers and theories?

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. 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 doublecone algebras. In the vacuum state of a conformal theory, the doublecone modular automorphism group acts as Hislop-Longo transformations (Hislop and Longo, 1982). Martinetti and Rovelli (2003) use this result to calculate the corresponding relationship between thermal time and proper time for a uniformly accelerating mortal observer:

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

where L is the observer's lifetime. (The relationship is more complicated in this case due to the fact that proper time is bounded while modular time is unbounded.) For most of the observer's lifespan, s is an approximately constant function of τ , allowing the Unruh temperature to again be interpreted as the local ratio between thermal and proper time.

This is the best we can hope for. Trebels (1997) proves that arbitrary doublecone automorphisms act as local dynamics, only if they act as scaled Hislop-Longo transformations.³ Of course, if nature is described by a non-conformal theory, then there is no guarantee that the doublecone modular automorphisms will have a suitable geometric interpretation. Saffary (2005) goes further, arguing that they will not have geometric significance in any theory with massive particles. The mathematical results backing this conjecture, however, are only partial.⁴

Attempting to generalize the TTH to cover non-uniform acceleration and non-vacuum states generates further difficulties. Work on the Unruh effect for non-uniformly accelerating observers (e.g., Jian-yang et al. 1995), indicates that such observers feel an acceleration-dependent thermal bath, reflecting the shifting ratio between constant thermal time and acceleration-dependent proper time. The TTH must explain the phenomenological experience of the observer who will presumably age according to her proper time, not the background thermal time flow. On top of this, if the global state is not a vacuum state, then it is not clear that the wedge modular automorphisms will carry a dynamical interpretation at all. The Radon-Nikodym theorem ensures that the action of the modular automorphism group uniquely determines the generating state. If ϕ, ψ are two (faithful, normal) states on a von Neumann algebra \mathfrak{M} , then the associated modular automorphism groups $\sigma_\phi^t, \sigma_\psi^t$ differ by a non-trivial inner automorphism, $\sigma_\phi^t(A) = U\sigma_\psi^t(A)U^*$, for all $A \in \mathfrak{M}, t \in \mathbb{R}$, so the general wedge dynamics will not be simple rescalings of the vacuum case.

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 conformal generalization. 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

³Formally, Trebels requires that local dynamics be continuous 1-parameter groups of automorphisms of the doublecone algebra that preserve subalgebra localization as well as spacelike and timelike relations between interior points. For a detailed discussion of Trebels's results, see Borchers (2000), §3.4.

⁴In the massless case, the modular generators are ordinary differential operators, δ_0 , of order 1. In the massive case, it has been conjectured that the modular generators are pseudo-differential operators $\delta_m = \delta_0 + \delta_r$, where the leading term is given by the massless generator δ_0 and δ_r is a pseudo-differential operators of order < 1 . This second term is thought to give rise to non-local action without geometric interpretation.

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 (e.g., 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.

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 . This justifies the clock hypothesis by linking the metrical properties of spacetime to the readings of tiny lightclocks. 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.

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 double-cone modular operators act geometrically like wedge operators implementing proper time translations along the observer's worldline. It is unlikely that the physics at this scale would 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 automorphisms associated with each (faithful, normal) state of a von Neumann algebra are connected by

inner automorphisms, they all project down onto the same 1-parameter group of outer automorphisms the algebra. 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, however, under what circumstances the outer flow acts in suitably geometric fashion to be interpretable as local dynamics, so it remains to be seen whether or not this is a viable option. The move does have immediate consequences for the global dynamics, however. Since the global algebra is expected to be type I, all modular automorphisms will be inner. As a result the canonical group of outer automorphisms is trivial. At a global level, there is no passage of time. At the local level, time emerges as a consequence of our ignorance of the global state.

4 The Classical Limit

The classical limit presents a different kind of challenge. Conceptually, 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 (i.e., $\phi(AB) = \phi(BA)$), and consequently every associated modular automorphism group acts as the identity, trivializing the thermal time flow. Does the TTH have a classical counterpart, or is quantum mechanics required to save time in a generally covariant setting?

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

$$\int_{\Gamma} dx \rho(x) = 1 \tag{7}$$

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

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

With respect to the corresponding Hamiltonian vector field, the evolution of

an arbitrary classical observable, $f \in C^\infty(\Gamma)$, is given by

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

and $\rho = \exp(-H_\rho)$. With respect to the Poisson bracket structure, the classical algebra of observables is non-abelian. Gallavotti and Pulvirenti (1976) use this non-abelian structure to define an analogue of the KMS condition. Is this connection strong enough to support a version of the TTH in ordinary general relativity? Or does it only serve to aid us in understanding how the thermal time variable behaves in the transition from quantum theory to classical physics?

The difficulty lies in connecting the thermal time flow for an arbitrary statistical state to our ordinary conception of time. In the quantum case this link was provided by the Bisognano-Wichmann theorem, which does not have a classical analogue. The problem is magnified by the lack of a full understanding of statistical mechanics and thermodynamics in curved space-time. 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.⁵

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*:

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

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,

⁵See Rovelli and Smerlak (2011).

describing how classical observables generate Hamiltonian vector fields on statespace. 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 Poisson manifold, Γ , Weinstein defines a corresponding *modular vector field* ϕ_μ given by the operator $\phi_\mu : f \rightarrow \text{div}_\mu H_f$ where H_f is the Hamiltonian vector field associated with a classical observable, $f \in C^\infty(\Gamma)$. The antisymmetry of the Poisson bracket entails that the operator ϕ_μ is a vector field on Γ . Weinstein proposes ϕ_μ as the classical analogue of the modular automorphism group. It characterizes the extent to which the Hamiltonian vector fields are divergence free (with respect to the density μ), vanishing iff all Hamiltonian vector fields are divergence free.

We can connect Weinstein's classical modular theory to the TTH. If Γ is a symplectic manifold and we let μ be the density associated with the canonical Liouville volume form, then $\phi_\mu(f) = 0$ for all observables. This reflects the conservation of energy by Hamiltonian flows in symplectic dynamical systems. Given any statistical state, however, we can define an associated density which leads to a nontrivial modular vector field. For any positive function, h , we have

$$\phi_{h\mu} = \phi_\mu + H_{-\ln h} = H_{-\ln h}. \quad (11)$$

Therefore any statistical state, ρ , defines a modular vector field equivalent to the Hamiltonian vector field $H_{-\ln \rho}$ associated with the density $e^{-\ln \rho} \mu$. We immediately recognize $-\ln \rho$ as the thermal Hamiltonian postulated by Connes and Rovelli. Clearly, $e^{is \ln \rho} \rho e^{-is \ln \rho} = \rho$, thus the state is invariant with respect to the flow of $H_{-\ln \rho}$. Additionally, it can be shown that ρ satisfies the KMS condition with respect to these dynamics, hence, from the perspective of the associated time flow ρ resembles an invariant equilibrium state just as in the quantum case.

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). There are, however, several deeper conceptual problems looming in the background which pose a more serious challenge to the viability of the hypothesis. Here, we will discuss two of the most pressing.

The first, which we will call the *generality problem*, draws upon the preceding discussion of the classical limit. While mathematically speaking, Weinstein’s modular vector field gives us a method for selecting a canonical thermal time flow in a classical theory, physical speaking, there is no reason why we should view the corresponding thermal time as physical time. As we have seen, any statistical state determines thermal dynamics according to which it is a KMS state, however, if ρ is a non-equilibrium state, the resultant thermal time flow does not align with our ordinary conception of time. By the lights of thermal time, a cube of ice in a cup of hot coffee is an invariant equilibrium state! The same problem arises in the quantum domain — only for states which are true equilibrium states will the thermal time correspond to physical time.

It appears inevitable that the TTH will have to be tempered. Rather than letting any state determine a corresponding flow of thermal time, only certain reference states should be permitted. Apart from the problem of providing an intrinsic, non-dynamical characterization of such states, if a system is not in one of these, it is hard to envision how a counterfactual state of affairs can determine the actual flow of time.⁶ This might provide more reasons for the defender of the TTH to explore the state independent, outer modular flow. Alternatively, she could try to argue that local non-equilibrium behavior can be viewed as small fluctuations from some background state. On this approach, the local flow of time in my office according to which the ice

⁶A closely related worry, what we might call the *background-dependence problem*, has been voiced by Earman (2011) and Ruetsche (2014). Their concern is that we can only identify modular automorphisms as dynamics because we already have a rich spatiotemporal geometry 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). This is exacerbated if the TTH is modified in response to the generality problem. Unless the modular automorphism group can always be viewed dynamically, the defender of the TTH will be hard pressed to find constraints capable of separating the dynamical cases from the non-dynamical cases which are independent of all background temporal structure.

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 (the entire universe perhaps). Rovelli (1993) hints in this direction, calculating that in a Friedman-Robertson-Walker universe, the thermal time induced by the equilibrium state of the cosmic microwave background will be proportional 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 localized, mortal observers like us. It would be truly remarkable to discover that our faculties of perception are sensitive to the thermal features of the CMB.

The second problem is the *gauge problem*. 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. 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. We lack a gauge invariant picture of generally covariant theories, and the TTH, at least in its present form, does not provide one.

Can a revised TTH give us the explanatory tools needed to understand the flow of time without reference to partial observables, or, does the entire framework of timeless mechanics require us to revise our conception of how ontology, explanation, and gauge symmetries are related?⁷ Whether or not

⁷Drifting in the latter direction, Rovelli (2014) suggests that gauge-dependent quantities are more than just mathematical redundancies, “they describe handles through 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.”

quantum thermodynamics can save time may rest on the solutions to these new incarnations of vexingly familiar philosophical problems.

References

- Alfsen, E. and F. Shultz (1998). Orientation in operator algebras. *Proceedings of the National Academy of Sciences, USA* 95, 6596–6601.
- Borchers, H. J. (2000). On revolutionizing quantum field theory with Tomita’s modular theory. *Journal of Mathematical Physics* 41(6), 3604–3673.
- 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.
- 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.
- 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.
- Hislop, P. D. and R. Longo (1982). Modular structure of the local algebras associated with a free massless scalar field theory. *Communications in Mathematical Physics* 84, 71.

- Jian-yang, Z., B. Aidong, and Z. Zheng (1995). Rindler effect for a nonuniformly accelerating observer. *International Journal of Theoretical Physics* 34, 2049–2059.
- 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.
- 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. (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.
- Trebels, S. (1997). *Über die Geometrische Wirkung Modularer Automorphismen*. Ph. D. thesis, Göttingen.
- Weinstein, A. (1997). The modular automorphism group of a Poisson manifold. *Journal of Geometry and Physics* 23, 379–394.