A Matter of Degree: Putting Unitary Inequivalence to Work

Laura Ruetsche†‡

If a classical system has infinitely many degrees of freedom, its Hamiltonian quantization need not be unique up to unitary equivalence. I sketch different approaches (Hilbert space and algebraic) to understanding the content of quantum theories in light of this non-uniqueness, and suggest that neither approach suffices to support explanatory aspirations encountered in the thermodynamic limit of quantum statistical mechanics.

1. Introduction. A characteristic, and provocative, feature of quantum field theory (QFT) is the availability of unitarily inequivalent Hilbert space representations of its canonical commutation relations (CCRs). Under the reasonable and historically entrenched assumption that unitary equivalence is a necessary condition for the physical equivalence of Hilbert space quantizations, this availability implies that there are myriad physically inequivalent quantizations of any classical field theory. I aim here to explore this dramatic non-uniqueness, and its implications for our understanding of the manner in which theories delimit physical possibilities. Lending both form and interest to this investigation is the existence of a level of abstraction at which even unitarily inequivalent Hilbert space quantizations share a common structure. They are, each of them, a concrete realization of an abstract algebraic structure—the structure of a C^* algebra called the *Weyl algebra,* and based upon the CCRs.

Each Hilbert space quantization is also, of course, a lot else, and it is by differing in features additional to their realization of the Weyl algebraic structure that quantizations can fail to be unitarily equivalent. Surveying

†To contact the author, write to Department of Philosophy, 1001 Cathedral of Learning, University of Pittsburgh, Pittsburgh, PA 15260; e-mail: ruetsche@pitt.edu.

‡For comments on earlier drafts, I am obliged to Gordon Belot, Jeremy Butterfield, Rob Clifton, and John Earman.

Philosophy of Science, 70 (December 2003) pp. 1329–1342. 0031-8248/2003/7005-0039$10.00

the roiling mob of inequivalent quantizations from the lofty heights of algebraic abstraction, one might suppose, as an early advocate of the algebraic approach to QFT does, that "all the physical content of the theory is contained in the algebra itself; nothing of fundamental significance is added to a theory by its expression in a particular representation" (Robinson 1966, 488). It would dissolve the foundational questions posed by the availability of unitarily inequivalent representations to deprive differences between those representations of physical significance.

Such a dissolution comes at a cost. In the QFT context, only a proper subset of the bounded operators on a Hilbert space representation of the CCRs instantiate the Weyl algebraic structure. Locating the physics solely in the abstract algebra exposes the remaining bounded operators—operators *parochial* to the representation—as unphysical accretions, clinging to concrete realizations of that algebra. Among the "accretions" in a particular representation are most of its projection operators, including those in the spectrum of its total number operator. Locating the physics in, and only in, the abstract algebra could mean investing with physical significance fewer observables than either scientific practice or our favored approaches to interpreting quantum theories can bear.

These reservations should not trigger retreat to a reactionary *Hilbert space chauvinism,* which identifies physically relevant observables with the set of bounded self-adjoint operators on some particular Hilbert space, and physically possible states with the set of density matrices on that Hilbert space. For the Hilbert space chauvinist runs the risk of investing with physical significance fewer *states* than our favored scientific and interpretive practices can bear. One option foreclosed is that of using states from unitarily *in*equivalent representations in our accounts of the phenomenon.

QFT is not the only setting where unitarily inequivalent representations arise. Quantum Statistical Mechanics (QSM), in its thermodynamic limit, is replete with unitarily inequivalent representations of its fundamental systems, infinite collections of microentities whose physics is quantum. What's more, explanations envisioned in the thermodynamic limit promise to complete the schematic arguments against chauvinism just offered. Thus the thermodynamic limit of QSM provides a motivation and a model for tempering chauvinisms, both Hilbert space and algebraic, about the structure and physical content of quantum theories.

I aim in what follows to make the foregoing somewhat more precise. Section 2 frames issues raised by unitarily inequivalent representations more ornately than I have so far. It reviews relevant rudiments of both Hilbert space and algebraic approaches to quantum theories, and describes Hilbert space and algebraic chauvinisms in more detail. Sketching the use to which QSM can put unitarily inequivalent representations, and an algebraic framework which encompasses them, in its treatments of equilibrium and phase transitions, Section 3 attempts to discredit both chauvinisms. In their stead, Section 4 offers an understanding of the content of physical theories which allows physical possibility to be a matter of degree.

2. Unitary Equivalence and Its Breakdown. Von Neumann discerned in both Schrödinger's wave and Heisenberg's matrix mechanics the structure of a *Hilbert space theory,* that is, a theory which (i) identifies the state space of a physical system with the set $\rho(\mathcal{H})$ of all positive normalized trace-class operators on a separable Hilbert space $\mathcal{H}$; (ii) associates the physical magnitudes (or observables) pertaining to that system with the set $\mathfrak{B}_{sa}(\mathcal{H})$ of all bounded, self-adjoint operators acting on $\mathcal{H}$; and (iii) where $\hat{\rho}$ is a state and $\hat{A}$ an observable, assigns $\hat{A}$ the expectation value $Tr(\hat{\rho}\hat{A})$ in the state $\hat{\rho}$.

One way to obtain such a structure is to successfully quantize a classical theory. A standard scheme for quantizing a theory cast in Hamiltonian form is to promote its canonical position and momentum observables to symmetric operators $(\hat{q}_i, \hat{p}_i)$ acting on a separable Hilbert space $\mathcal{H}$ and satisfying CCRs answering to the classical Poisson bracket.

Call any set of Hilbert space operators that does the trick a *representation* of the CCRs. According to a theorem announced in 1930 by Stone and proven the next year by von Neumann, if $(\mathcal{H}, \{\hat{O}_i\})$ and $(\mathcal{H}', \{\hat{O}_i'\})$ are both irreducible representations of the CCRs for finitely many degrees of freedom,[1] then $(\mathcal{H}, \{\hat{O}_i\})$ and $(\mathcal{H}', \{\hat{O}_i'\})$ are *unitarily equivalent*, that is, there exists a one-to-one, linear, norm-preserving transformation ("unitary map") $U : \mathcal{H} \rightarrow \mathcal{H}'$ such that $U^{-1}\hat{O}_i' U = \hat{O}_i$ for all values of i. It follows not only that Heisenberg's matrix realization of the CCRs for n degrees of freedom is unitarily equivalent to Schrödinger's representation in terms of differential operators, but also that *any* Hilbert space representation of the CCRs for n degrees of freedom is equivalent to the Schrödinger representation.

The Stone-von Neumann theorem was widely received as proof that the physical theory arising from the quantization of an n-dimensional classical theory is essentially unique. A number of non-trivial assumptions about the nature and content of quantum theories underlies this reception. To

1. A representation of the CCRs is *irreducible* if and only if the only subspaces of $\mathcal{H}$ invariant under the action of all operators in the representation are the zero subspace and $\mathcal{H}$ itself. There are further technical assumptions; for an introduction, see Wald 1994, Ch. 2.2. (To state the obvious: I don't aim here at a presentation comprehensive in technical detail, and so will refer interested readers to more rigorous discussions. To avoid expository clutter, some peripheral technical notions will be defined in footnotes rather than in the text.)

reconstruct these assumptions, start with the idea that *the content of a physical theory is the set* Ω *of worlds possible according to the theory.* Annex to this an assumption about how a statistical theory characterizes a physical possibility: *A physical possibility* $\omega \in \Omega$ *is an assignment of expectation values to a set* $\mathcal{A}$ *of physical magnitudes*. Then one may denote the content of a physical theory by the pair $(\Omega, \mathcal{A})$, where Ω is its set of possibilities, that is, maps $\omega : \mathcal{A} \to \mathbb{R}$ from its set $\mathcal{A}$ of physical magnitudes to their expectation values. Theories will be physically equivalent exactly when a suitable isomorphism obtains between the sets of possibilities they recognize—more precisely, when they satisfy a content coincidence criterion for physical equivalence, which Clifton and Halvorson 2001 articulate as follows:

> $(\Omega, \mathcal{A})$ and $(\Omega', \mathcal{A}')$ are physically equivalent if and only if there exist bijections $i_s : \Omega \to \Omega'$ and $i_o : \mathcal{A} \to \mathcal{A}'$ such that for all $\omega \in \Omega$ and for all $A \in \mathcal{A}$, $\omega(A) = [i_s(\omega)](i_o(A))$.

To complete the case that quantum theories are physically equivalent if and only if they're unitarily equivalent, one now need only assume that quantum theories are Hilbert space theories, in the sense of satisfying conditions (i)–(iii) announced in the first paragraph of this section. Then it follows that quantum theories $(\rho(\mathcal{H}), \mathfrak{B}_{sa}(\mathcal{H}))$ and $(\rho(\mathcal{H}'), \mathfrak{B}_{sa}(\mathcal{H}'))$ are physically equivalent if and only if they're unitarily equivalent, in which case the unitary map furnishes both the bijection of possibilities from $\rho(\mathcal{H})$ to $\rho(\mathcal{H}')$ and the bijection of magnitudes from $\mathfrak{B}_{sa}(\mathcal{H})$ to $\mathfrak{B}_{sa}(\mathcal{H}')$, by mapping an operator X on $\mathcal{H}$ to an operator $X' = U^{-1}XU$ on $\mathcal{H}'$. (For a proof, see Bratteli and Robinson 1987, Theorem 2.3.16.)

Filtered through these presuppositions, the Stone-von Neumann theorem issues remarkable reassurance that the Hamiltonian quantization of a finite dimensional classical theory issues a unique quantum theory. Physicists embracing the assumptions catalogued here can quantize classical theories with confidence that the upshot will be an unambiguous and coherent quantum theory.

Unless, that is, the theory they're quantizing falls outside the scope of the Stone-von Neumann theorem. The uniqueness result holds only for quantizations of classical theories whose configuration spaces are finite in dimension. Elsewhere, uniqueness breaks down, and breaks down dramatically. Where X is a classical *field* theory, continuously many unitarily inequivalent quantizations can vie for the title "quantization of X." Supposing that unitary equivalence is criterial for physical equivalence, one must also suppose that at most one unitary equivalence class of quantizations can hold the title. This being supposed, the availability of unitarily inequivalent representations renders "the quantization of X" ambiguous at best, if not incoherent.

One reaction to this state of affairs is to disambiguate—to specify that unitary equivalence class of representations in which resides the content of the theory. This reaction entails dismissing representations unitarily inequivalent to the privileged one—and operators parochial to those representations—as without physical significance. Pledging allegiance to the idea that the space of quantum theoretic possibilities is given in terms of a fixed Hilbert space by $(\rho(\mathcal{H}), \mathfrak{B}_{sa}(\mathcal{H}))$, this is the reaction of the Hilbert space chauvinist.

Investigations of the abstract structure of standard realizations of quantum field theoretic CCRs conducted in the 50s and 60s inspire another reaction. These investigations revealed that each concrete Hilbert space representation of the CCRs gives rise to an abstract algebra C^* algebra, the *Weyl algebra*, which is representation-independent. (For more on algebraic notions introduced in this section, see Wald 1994.) To frame the interpretive stance which rests on this representation-independence, I will sketch the rudiments of an algebraic approach to quantum theories. The algebraic approach identifies quantum observables with self-adjoint elements of a C^* algebra $\mathcal{A}$ (i.e., elements A such that $A^* = A$). This algebra can be abstract, like the Weyl algebra, or it can be an algebra of bounded operators on a fixed and concrete Hilbert space. Thus the algebraic approach generalizes the Hilbert space notion of observable.

It is with respect to observables in the more general sense of elements of an algebra that the algebraic approach constitutes its notion of state. An algebraic state ω on $\mathcal{A}$ is a linear functional $\omega : \mathcal{A} \to \mathbb{C}$ that is normed ($\omega(I) = 1$)) and positive (ω $(A^*A) \geq 0$ for all $A \in \mathcal{A}$). Hence $\omega(A)$ may be understood as the expectation value of an observable $A \in \mathcal{A}$. Where the algebra is realized as bounded operators on $\mathcal{H}$, the set of countably additive algebraic states stands, via the trace prescription, in one-to-one correspondence with the set $\rho(\mathcal{H})$ of density operators on $\mathcal{H}$. But the general notion of an algebraic state does not require a Hilbert space middleman.

The set Ω of states on a C^* algebra $\mathcal{A}$ is convex. Its extremal elements—that is, states ω which cannot be expressed as non-trivial convex combinations of other states—are pure states; all other states are mixed.

Associations can be drawn between algebraic and Hilbert space frameworks. A *Hilbert space representation* of an abstract algebra $\mathcal{A}$ is a structure-preserving map $\pi : \mathcal{A} \to \mathfrak{B}(\mathcal{H})$, from elements of $\mathcal{A}$ to the set of bounded linear operators on a Hilbert space $\mathcal{H}$. That even abstract algebras admit concrete Hilbert space representations enables us to connect Hilbert space and algebraic notion of states. A state $\hat{\rho}$ in a Hilbert space $\mathcal{H}$ carrying a representation $\pi : \mathcal{A} \to \mathfrak{B}(\mathcal{H})$ of an algebra $\mathcal{A}$ naturally gives rise to the algebraic state $\omega(A) = Tr(\hat{\rho}\pi(A))$ for all $A \in \mathcal{A}$.

We can also move in the other direction, from an algebraic state to its realization on a concrete Hilbert space. A state ω over a C^* algebra $\mathcal{A}$ can be recast as state in a Hilbert space bearing a faithful[2] representation of that algebra. That is, for such a state, there exists a Hilbert space $\mathcal{H}_\omega$, a faithful representation $\pi_\omega : \mathcal{A} \rightarrow \mathfrak{B}(\mathcal{H}_\omega)$, and a cyclic[3] vector $|\Psi_\omega\rangle \in \mathcal{H}_\omega$ such that $\omega(A) = \langle\Psi_\omega|\pi_\omega(A)|\Psi_\omega\rangle$ for all $A \in \mathcal{A}$. Called the GNS representation of the state (for Gel'fand, Naimark, and Segal, who showed how to construct it), the triple $(\mathcal{H}_\omega, \pi_\omega, |\Psi_\omega\rangle)$ is unique up to unitary equivalence. An algebraic state ω is pure if and only if its GNS representation π_ω is irreducible. Mixed algebraic states give rise to reducible GNS representations.

The following facts and locutions will be called into service down the road. The *folium* of an algebraic state ω is the set of all algebraic states which may be expressed as density matrices on ω's GNS representation. Suppose that algebraic states ω and ω' are pure. Then either they give rise to unitarily equivalent GNS representations, or they do not. In the first case, their folia coincide. In the second, their folia are *disjoint*, that is, no algebraic state expressible as a density matrix on $(\mathcal{H}_\omega, \pi_\omega)$ is expressible as a density matrix on $(\mathcal{H}_{\omega'}, \pi_{\omega'})$, and vice versa. A fact that will come to the fore in Section 3 is that mixed algebraic states can be convex combinations of disjoint algebraic states: consider $\omega = \lambda\omega_1 + (1 - \lambda)\omega_2$, with ω_1 and ω_2 disjoint.

Setting this framework for algebraic quantum theory alongside the result that the Weyl algebra is representation-independent one might think that where there's a Weyl (algebra), there's a way (to do QFT). The position I'll call *algebraic chauvinism* denies that quantum theories are essentially Hilbert space theories. The algebraic chauvinist identifies all physical magnitudes pertaining to a system with self-adjoint elements of its Weyl algebra, and takes the complete set of states possible for this system to be given by normed positive linear functionals ω over this abstract algebra. For the chauvinist, "the important thing here is that the observables form some algebra, and not the representation Hilbert space on which they act" (Segal 1967, 128). Withholding significance from representation-dependent structures, algebraic chauvinists have the luxury of greeting unitarily inequivalence with a yawn.

3. Unitarily Inequivalent Representations in QSM. This section sketches some uses to which QSM would put unitarily inequivalent representations, uses which, I suggest, should give both chauvinists pause.

2. A representation π of $\mathcal{A}$ is *faithful* if and only if $\pi(A) = 0 \Rightarrow A = 0$.

3. $|\Psi\rangle$ is cyclic for $\pi_\omega(\mathcal{A})$ means $\{\pi_\omega(\mathcal{A})|\Psi\rangle\}$ is dense in $\mathcal{H}$.

(Quantum) statistical mechanics aims at a (quantum) microphysical underpinning of bulk properties associated with macrosystems—their temperature, pressure, entropy, and the like—an underpinning which stands in some suitable explanatory relationship to thermodynamic laws those macro-properties obey. Articulating canons of suitability, and assessing putative statistical mechanical explanations against those canons, has been a mainstay of work on the foundations of statistical mechanics. Until my concluding anticlimactic postscript, I will bracket questions raised by such work, in order to explore certain well-entrenched statistical mechanical explanatory aspirations, and the theoretical/interpretational structures which sustain them. I will focus in particular on explanatory aspirations pursued in the *thermodynamic limit* of QSM, i.e. the limit as the number N of microsystems and the volume V they occupy goes to infinity, while their density $\frac{N}{V}$ remains finite. Because the thermodynamic limit for QSM concerns the quantum physics of infinite collections of particles, there the spectre of unitarily inequivalent representations rears its head.

Short of the thermodynamic limit, and in the setting of concrete Hilbert spaces, the *Gibbs state* equips QSM with a notion of equilibrium. The Gibbs state of a system with Hamiltonian $\hat{H}$ at inverse temperature $\beta = \frac{1}{kT}$ is the density matrix

$$\hat{\rho} = \exp(-\beta\hat{\mathrm{H}})/\mathrm{Tr}[\exp(-\beta\hat{\mathrm{H}})] \qquad (1)$$

For realistic, finite quantum systems the Gibbs state is well-defined and unique (Ruelle 1969). If, however, the spectrum of $\hat{H}$ fails to be pure discrete, or if we are working in an abstract algebraic setting, (1) fails to be well defined.

Now suppose that we aspire to construct a quantum statistical account of *phase transitions*. Then we might have reason to seek a notion of equilibrium suited to these more general settings. For the apparent macroscopic explanandum is the existence, at certain temperatures, of multiple thermodynamic phases. The explanatory aspirations I'll recount here rest on the idea that a statistical account of phase transitions requires the existence, at these critical temperatures, of *multiple distinct* equilibrium states, answering to different thermodynamic phases. Sewell explains how, short of the thermodynamic limit, the very uniqueness of the Gibbs state upsets this explanatory applecart:[4]

> The traditional form of statistical thermodynamics for large but finite systems . . . cannot accommodate different phases of a system (e.g.

4. Another reason for going to the thermodynamic limit, which I do not consider here, is that it is only in the thermodynamic limit that discontinuities in thermodynamic functions, which discontinuities are characteristic of phase transitions, occur.

> liquid and vapor) under the same thermodynamic conditions, since the Gibbs ensemble representing the equilibrium state of a *finite* system is uniquely determined by the prevailing macroscopic constraints: thus, if the volume, temperature, and mass are controlled to take specific values, then the resultant ensemble is the canonical one. (1986, 47)

Explanatory hopes are revived in the thermodynamic limit by using the *KMS condition* to explicate a notion of equilibrium more general than that afforded by the Gibbs state. (A naive introduction to KMS states follows; for an authoritative treatment, see Bratteli and Robinson 1997, §5.3.) A C^* dynamical system $(\mathcal{A},\alpha_t)$ consists of a C^* algebra $\mathcal{A}$ whose self-adjoint elements correspond to physical magnitudes, and a one (real) parameter group α_t of automorphisms on $\mathcal{A}$—that is, maps from $\mathcal{A}$ to itself which preserve $\mathcal{A}$'s algebraic structure—which encodes dynamics. That is, for all $A \in \mathcal{A}$, $\alpha_t(A)$ represents its evolution through a time t. In a Hilbert space quantum theory, $\mathcal{A}$ is given by an algebra of bounded observables on a Hilbert space, and α_t is implemented by a family $\hat{U}_t = e^{-i\hat{H}t}$ of unitary operators generated by the Hamiltonian $\hat{H}$ of the system: $\alpha_t(\hat{A}) = \hat{U}_t\hat{A}\hat{U}_t^*$.

In terms of such a Hilbert space realization of a C^* dynamical system, the Gibbs state $\hat{\rho}$, where it is well-defined, formally satisfies

$$\omega[\mathrm{A}\alpha_{\mathrm{i}\beta}(\mathrm{B})] = \omega(\mathrm{BA}) \text{for all } \mathrm{A}, \mathrm{B} \in \mathcal{A} \tag{2}$$

(here $\omega(x) = Tr(\hat{\rho}x)$ for all $x \in \mathfrak{B}(\mathcal{H})$). But formulated in general C^* algebraic terms, (eq. 2) can apply as well to states and observables abstractly conceived. To extrapolate the notion of equilibrium beyond circumstances where the Gibbs state (eq. 1) is well-defined, make the KMS condition (2) criterial for equilibrium. Hence: ω is a *KMS state* with respect to the automorphism group α_t at inverse temperature β (an (α_t, β)-KMS state, for short) if and only if (2) holds for all A, B in a dense subalgebra of $\mathcal{A}$.

If $(\mathcal{A}, \alpha_t)$ admits a standard Gibbs state at inverse temperature β, the (α_t, β)-KMS state is unique and coincides with that Gibbs state (Bratteli and Robinson 1997, Ex. 5.3.31). KMS states moreover exhibit a number of stability features, including invariance under the action of the dynamical group α_t, putatively characteristic of equilibrium states. For such reasons, the KMS condition is generally regarded to be a suitable criterion for equilibrium. So explicated, the notion applies to systems admitting no Gibbs states—including infinite quantum systems at the thermodynamic limit.

Now consider a C^* dynamical system $(\mathcal{A}, \alpha_t)$. For $\beta \in \mathbb{R}$, let K_β denote the set of (α_t, β)-KMS states. Salient results about the structure of the sets K_β include (Bratteli and Robinson 1997, Theorem 5.3.30):

1. K_β is convex;
2. $\omega \in K_\beta$ is extremal (i.e., ω can't be expressed as a convex combination of distinct elements of K_β) if and only if it's a *factor* state (i.e., one for which the intersection of $\pi_\omega(\mathcal{A})$ and its commutant contains only multiples of the identity);
3. Where ω_1 and ω_2 are extremal elements of K_β, either they're equal or disjoint.

It follows from (2) that if the set of (α_t, β)-KMS states has only one element, then that state is a factor state. Factor states, examples of which include the equilibrium states of ideal fermi gases, can often be characterized by the absence of long-range correlations, and of large fluctuations for space-averaged observables. Typical of "pure" thermodynamic phases, these absences encourage the identification of factor states with those phases (for more encouragement, see Sewell 1986, § 4.4, or Emch and Knops 1970).

Now consider $\omega_1 \in K_{\beta_1}$ and $\omega_2 \in K_{\beta_2}$. Under a technical assumption that holds generally at the thermodynamic limit,[5] if $\beta_1 \neq \beta_2$, then ω_1 and ω_2 are disjoint (Bratteli and Robinson 1997, theorem 5.3.35). That is, for an infinite quantum statistical system, there is no single concrete Hilbert space on which its equilibrium states at different temperatures can be represented as density matrices.

The position of the Hilbert space chauvinist, viewed in the light of this result, looks unreasonable. Maintaining that all physical possibilities reside in a single folium, the chauvinist reckons states outside the favored folium to be physically impossible. But there are systems for which this amounts to insisting that at most one equilibrium temperature is physically possible. The Hilbert space chauvinist cannot allow that it's in some sense physically possible for such systems to reach equilibrium at different temperatures. While not inconsistent, this consequence offends modal intuitions.

Setting our sights on the explanation of phase transitions only makes Hilbert space chauvinism look worse. Recall that for finite systems admitting Gibbs states, the equilibrium (KMS) state at temperature β with respect to an automorphism group α_t is unique. But in the general setting of the thermodynamic limit of QSM, there can be automorphism groups α_t and inverse temperatures β such that there are a plurality of (α_t, β) KMS states. Every ω in such a set K_β can be represented as a *unique* convex combination of extremal elements of K_β, which extremal elements are pairwise disjoint (Bratteli and Robinson 1997, Theorem 5.3.30).

5. Where $\mathcal{U}'$ denotes the commutant of $\mathcal{U}$, the assumption is that the von Neumann algebras $\pi_{\omega_1}(\mathcal{A})''$ and $\pi_{\omega_2}(\mathcal{A})''$ are Type III; for a sketch of why it holds generally at the thermodynamic limit, see Emch 1984, 448–450.

This makes available the following template for a quantum statistical analysis of phase transitions. Phase transitions occur at those inverse temperatures β for which the set K_β of (α_t, β) KMS states is not a singleton set and in those states $\omega \in K_\beta$ which are not extremal. Such ω are convex combinations of extremal states ω_i. Each extremal state in this decomposition corresponds to a pure thermodynamic phase, different states to different phases. Thus the decomposition corresponds to the separation of a system at equilibrium into pure thermodynamic phases, and a system in ω at β exhibits phase transitions (see Sewell 1986, ch. 4).

The analysis of phase transitions just sketched takes disjoint algebraic states ω_i to co-exist in the form of different phases present at a phase transition. Implying that what's *actual* can on its own correspond to multiple, distinct folia, this explanation is incompatible with Hilbert space chauvinism, which limits the space of physical possibilities to a single folium.

Pursuing explanatory aspirations in the thermodynamic limit requires extending physical possibility beyond the lone folium to which a Hilbert space chauvinist would confine it. It does not follow that algebraic chauvinism is ideally supportive of the account of phase transitions just sketched. The algebraic chauvinist's catechism is that moving to a concrete representation adds no *physical* content to a theory couched in terms of an abstract algebra. The foregoing might tempt one to protest that concrete representations bear crucial physical content, corresponding as they do to the phase and the temperature of a system at equilibrium. This on its own needn't trouble the chauvinist, provided that she can understand the admittedly physical differences between unitarily inequivalent representations—differences of phase and of temperature—in purely algebraic terms. If she could, for instance, summon from her algebra a self-adjoint element T such that for any (α_t, β) KMS state ω, $\omega(T) = \frac{1}{\beta k}$, then purely algebraic resources would suffice for the temperature discriminations her critic assigns to concrete representations.

Alas, algebraic resources do not, on their own, supply the chauvinist with a temperature observable. To indicate why, we must articulate the algebraic approach to QSM in a bit more detail (Primas 1983, §4.3 offers a sketch; see also Kronz and Lupher 2001). That approach associates with each bounded region V of $\mathbb{R}^3$ (i.e., three dimensional physical space) an algebra $\mathcal{A}(V)$. Call elements of such algebras *strictly local observables*. From these algebras is constructed a C^* algebra, the *quasi-local algebra* $\mathcal{A}$, roughly as follows: take $\bigcup_{V \in \mathbb{R}^3} \mathcal{A}(V)$, then close in the topology furnished by the C^* algebraic norm. According to the algebraic chauvinist, it is in this abstract quasi-local algebra that all physically relevant observables reside. Now the rub is that classical thermodynamic observables, including temperature, are absent from this quasi-local algebra. That is, they are not observables in terms of which the algebraic chauvinist can distinguish between states.

The accounts of equilibrium and phase transitions just sketched extend physical possibility to unitarily inequivalent representations. Therein lies their incompatibility with Hilbert space chauvinism. Those same accounts distinguish physically between those representations on the basis of observables without correlate in the abstract algebra. Therein lies their incompatibility with algebraic chauvinism.

Now, algebraic approaches unfettered by chauvinism can bring classical thermodynamic observables on board. Here's (again, roughly) how. For every region V define the algebra $\mathcal{A}^{\perp}(V)$ of *quasi-local observables outside* V by taking the norm closure of $\{A : A \in \mathcal{A}(V'),\ V' \cap V = \emptyset\}$. Given a Hilbert space representation π of the quasi-local algebra $\mathcal{A}$, one can construct a von Neumann algebra $\mathcal{V}_{\pi}^{\perp}(V)$ by taking the closure of $\pi(\mathcal{A}^{\perp}(V))$ in the weak operator topology of the representation's Hilbert space. The *von Neumann algebra 'at* ∞' $\mathcal{V}_{\pi}^{\infty}$ is defined by $\bigcap_{V \in \mathbb{R}^3} \mathcal{V}_{\pi}^{\perp}(V)$ (see Bratteli and Robinson 1987, 119–122). It is in this construction, obtained only by way of a representation π, that one encounters classical thermodynamic observables.

4. Conclusion: Coalescing Content. How can we construe the content of quantum theories in a way that accommodates explanatory maneuvers encountered in the thermodynamic limit of QSM? I've just suggested that neither chauvinism will do. But perhaps they are not the only options open to us. Their opposition notwithstanding, algebraic and Hilbert space chauvinism share an assumption about how to interpret a physical theory. The shared assumption is that a physical theory's content is to be specified by simply sorting logical possibilities into one of two disjoint and exhaustive categories: the physically possible and the physically impossible—as though a physical theory ran a modal toggle with no intermediate settings. The Hilbert space chauvinist uses a Hilbert space structure of observables to do the simple sort; the algebraic chauvinist uses the abstract algebraic structure.

Mathematical physicists discussing algebras and their representations might be taken to suggest a different take on how physical theories pick out possibilities. The remainder of this section aims not at a fully satisfying explication of the suggestion, but at a partial, and admittedly impressionistic, development of it.

Kadison makes the suggestion this way:

> Mathematically, a representation [of an abstract algebra] distinguishes a certain "coherent" family of states from among [the full set of

> algebraic states], and at the same time, in effect, ''coalesces'' some of the algebraic structure. (1965, 186)

The distinguished family is the folium of states expressible as density matrices on the Hilbert space of the representation; the coalesced structure includes observables parochial to that representation and accessible through constructions—e.g. the von Neumann algebra at ∞—based on that representation.

Kadison takes concrete representations seriously as repositories of physical content without assuming that *all* physically relevant states must reside in a single folium. He thereby suggests how to chart a course (a course I'll call, for reasons which will become apparent, the *Swiss army approach*) between algebraic and Hilbert space chauvinisms. The Swiss army approach has as its point of departure a refusal to specify the content of a physical theory in one fell swoop, cleaving states possible according to it from states impossible according to it. Rather, the Swiss army approach takes the specification of content to be (at least) a two-tiered affair, with a corresponding gradation in the sort of possibilities purveyed by the theory. The broadest sort of possibility picked out by a quantum theory is the space $\Omega_{\mathcal{A}}$ of algebraic states on the appropriate abstract algebra. Self-adjoint elements of the algebra correspond to the most basic physical magnitudes, those that belong to the theory automatically. This much of the theory's content can be specified, so to speak, a priori, before taking physical contingencies into account.

The next tier of physical content specification does take contingencies into account. From $\Omega_{\mathcal{A}}$ a narrower set of possibilities most relevant to the contingent empirical situation is distinguished, by appeal to features of that situation, for instance, equilibrium temperatures. Other algebraic states aren't impossible; they're simply possibilities more remote from the present application of the theory than these most relevant states.

This narrowing of possibilities expands the core constituency of relevant observables from $\mathcal{A}$ to include observables parochial to concrete GNS representations of states in the narrower set. Again, observables parochial to the GNS representations of states outside the most relevant set aren't *unphysical*; they're simply less relevant for the sorts of discriminations demanded by the application at hand. This is hardly to say that they couldn't be relevant to other applications.

We can also think of this two-tiered specification of content in terms of the *universal representation* of an algebra $\mathcal{A}$. This is the direct sum, over the set of algebraic states for $\mathcal{A}$, of their GNS representations. We could construe the theory's broadest set of physical observables in terms of this universal representation (see Kronz and Lupher 2001 for one version of this proposal, which they attribute to Muller-Herold; Rob Clifton has also offered a version of this proposal in conversation). At this stage of content specification, this vast host of physical observables is just sitting there, like blades folded up in a Swiss army knife. The next (coalescence) stage appeals to contingent features of the physical situation to focus on a small set of representations, which are summands in the universal representation. Observables parochial to those representations are extracted for application to the situation at hand. Thus coalescence is something like opening the Swiss army knife to the appropriate blade or blades, once you've figured out what you're supposed to do with it.

Though observables thus coalesced are less fundamental than those appearing in $\mathcal{A}$, there can be call to draft them. We've just seen how the coalesced von Neumann algebra at ∞ sustains quantum statistical explanatory aspirations. Coalesced observables can be pressed into other sorts of service. For example, a Hamiltonian parochial to a representation might serve as the generator of dynamics in the folium of that representation. (Emch and Knops's (1970) variation on the Ising model of ferromagnetism develops a dynamics of this sort.) The interplay of abstract and coalesced structure also figures in accounts which characterize spontaneous symmetry breaking in terms of symmetries of the algebraic structure which are not symmetries of coalesced structures. In the case of phase transitions, a non-extremal KMS state $\omega = \Sigma\lambda_i\omega_i$ might be invariant under symmetries (automorphisms) of the abstract algebra which fail to be unitarily implementable on representations coalesced around that states' extremal components ω_i.

My suggestions that the Swiss army approach admits and supports the foregoing applications are sketchy, and the Swiss army approach itself remains a metaphor. But if these ideas can be developed and defended, they would make plausible the thesis that the best way to make sense of what physicists do with quantum theories is to allow physical possibility to be a matter of degree.

I'll close by acknowledging an objection[6] to how I've proceeded. Steam rises from the surface of my coffee; passing to the thermodynamic limit to account for this, we attribute my coffee cup infinite volume. But the volume of my coffee cup is finite! So the objection is that I have rested interpretative conclusions on the consideration of a setting which is a hotbed of manifest falsehoods and extreme idealizations.

My shamefully curt reply to this objection is that I am not resting interpretative conclusions on artifacts of the idealization committed by the thermodynamic limit. I am not (for instance) claiming that, notwithstanding the appearances, steaming cups of coffee are infinite in volume. I am instead

6. John Earman, brandishing a copy of Callender (forthcoming), has urged this objection upon me.

resting interpretative conclusions on the very features of the idealization—in particular, the structure of equilibrium states it sustains—that enables it to account for the phenomena—the coexistence of phases at critical temperatures. Thus I am resting interpretive conclusions—conclusions about the manners in which theories represent—on those facets of the thermodynamic limit that appear to do representational work. But on this question, much more needs to be said, both by the prosecution and the defense.

REFERENCES

Bratteli, Ola, and Derek W. Robinson (1987), *Operator Algebras and Quantum Statistical Mechanics 1*, 2nd ed. Berlin: Springer-Verlag.

——— (1997), *Operator Algebras and Quantum Statistical Mechanics 2*, 2nd ed. Berlin: Springer-Verlag.

Callender, Craig (2001), "Taking Thermodynamics Too Seriously", *Studies in the History and Philosophy of Modern Physics* 32: 539–554.

Clifton, Rob, and Hans Halvorson (2001), "Are Rindler Quanta Real?", *British Journal for the Philosophy of Science* 52: 417–470.

Emch, Gerard G. (1984), *Mathematical and Conceptual Foundations of Twentieth Century Physics*. Amsterdam: North-Holland.

Emch, Gerard G., and H.J.F. Knops (1970), "Pure thermodynamic phases as extremal KMS states", *Journal of Mathematical Physics* 11: 3008–3018.

Kadison, Richard (1965), "Transformations of States in Operator Theory and Dynamics", *Topology* 3, Suppl. 2: 177–198.

Haag, Rudolf (1962), "The Mathematical Structure of the Bardeen-Cooper-Schrieffer Model", *Nuovo Cimento* 25: 287–299.

Kronz, Fred, and Tracy Lupher (2001), "Unitarily Inequivalent Representations in Algebraic Quantum Theory", pre-print.

Primas, Hans (1983), *Chemistry, Quantum Mechanics, and Reductionism*. New York: Springer-Verlag.

Robinson, Derek W. (1966), "Algebraic Aspects of Relativistic Quantum Field Theory", in M. Chretien and S. Deser (eds.), *Axiomatic Field Theory*. New York: Gordon and Breach.

Ruelle, David (1969), *Statistical Mechanics*. New York: W.A. Benjamin.

Segal, Irving (1967), "Representation of Canonical Commutation Relations", in Francois Lurcat (ed.), *Cargese Lectures in Theoretical Physics*. New York: Gordon and Breach.

Sewell, Geoffrey L. (1986), *Quantum Theory of Collective Phenomena*. Oxford: Oxford University Press.

Wald, Robert M. (1994), *Quantum Field Theory in Curved Spacetime and Black Hole Thermodynamics*. Chicago: University of Chicago Press.