An Alternative Interpretation of Statistical Mechanics

C.D. McCoy*†

26 September 2016

Abstract

Abstract: I propose an interpretation of classical statistical mechanics that centers on taking seriously the idea that probability measures are complete statistical mechanical states. I show how this leads naturally to the idea that the stochasticity of statistical mechanics is associated directly with the observables of the theory rather than with the microstates (as traditional accounts would have it). The usual assumption that microstates necessarily have representational significance in the theory is therefore dispensable, a consequence which reveals interesting possibilities for investigating inter-theoretic explanations of the foundational questions of statistical mechanics.
Word Count: 9753

1 Introduction

Commonly lamented in the literature on the foundations of statistical mechanics is the lack of a canonical formalism on which to base foundational discussions. Unlike in non-relativistic quantum mechanics, general relativity, or classical mechanics, "what we find in [statistical mechanics] is a plethora of different approaches and schools, each with its own programme and mathematical apparatus, none of which has a legitimate claim to be more fundamental than its competitors" (Frigg, 2008, 101). Because of this, "the philosophical foundations of thermodynamics and statistical mechanics can seem a bewildering labyrinth" writes Callender (2011, 83), and, indeed, it seems to many workers in the field that "consequently we have no choice but to dwell on its history" (Uffink, 2007, 923) when investigating the foundations of the theory.

While that history has indeed witnessed many approaches to the foundations of statistical mechanics, a common theme emerges in a considerable number of the debates, a theme emphasized throughout Sklar's panoptic survey of the field, *Physics and Chance*. There is often a tension between preserving some thermodynamic regularity or fact exactly, or else insuring that the theory describes individual microscopic systems (Sklar, 1993; Callender, 2001). Choosing the former option has usually led to the adoption of an "ensemble" concept of statistical mechanics, which evokes "the imaginative picture of an innumerably vast number of systems, all subject to the same macroscopic constraint and taking on the infinite variety of possible microscopic states compatible with those constraints" (Sklar, 1993, 159). Choosing the latter

*School of Philosophy, Psychology, and Language Sciences, University of Edinburgh, Edinburgh, UK. email: casey.mccoy@ed.ac.uk

†**Acknowledgements:** My thanks to Harvey Brown, Craig Callender, Richard Dawid, John Dougherty, Gijs Leegwater, F.A. Muller, Patricia Palacios, Anncy Thresher, David Wallace, Stefan Wintein, and Christian Wüthrich for helpful comments and conversations. I also wish to thank audiences at the Vienna Circle Institute and the Munich Center for Mathematical Philosophy, where some of this material was presented.

option, popularly called the "Boltzmannian" approach, leads one to give up the exact laws in favor of a statistical regularities, while preserving the idea that each individual system has an exact microstate.[1]

As many commentators have noted, the ensemble interpretation of statistical mechanics makes it unclear how statistical mechanics can explain the behavior of individual systems, which has led many philosophers in recent years to favor the Boltzmannian approach. There is, however, an alternative that has received scant attention from philosophers. In this paper I propose an interpretation of statistical mechanics that parts ways with the two horns of Sklar's dilemma. I claim that a coherent and illuminating way to understand statistical mechanics is to take seriously the idea that statistical mechanical states are probability measures. Indeed, this is to some extent assumed in both the ensemble and Boltzmannian interpretations, but I mean something distinct. I claim that statistical mechanical systems have individual states (as in the Boltzmannian interpretation), but that these states are *completely* represented in the theory by probability measures (as in the ensemble interpretation). It is then a natural step from this claim to the idea that the real, physical stochasticity of statistical mechanical systems is located in the behavior of macroscopic observables; that is to say, the observable properties of these systems are objectively random (at least from the point of view of classical statistical mechanics).

It follows from the stochasticity of the observables that one cannot consistently conceive of statistical mechanical systems in the Boltzmannian way, i.e. as possessing deterministically-evolving, microscopic, mechanical states in addition to the macroscopic state represented by a probability measure. This is not to say, I emphasize, that I am advocating an instrumentalist interpretation of the theory. There plausibly is always a microscopic state of any statistical mechanical system, one that is properly explicated in a more fundamental theory (like quantum mechanics). Indeed, a major motivation of this interpretation is the idea that the foundations of statistical mechanics is best explored *inter-theoretically* rather than *intra-theoretically*, especially by attending to the classical limit of quantum systems, as has been urged especially by Wallace (2001, 2015).

Thus, one of my principal contentions through advocating this alternative interpretation is that one should simply forgo interpreting microstates as *necessarily* representing an underlying particle ontology of classical statistical mechanics.[2] *That is not to say that some statistical mechanical systems appropriately reduce to classical systems of particles*. In some circumstances reasonable cases can be and have been made. My charge is that assuming that they all do necessarily is unwarranted by what we presently know about physics. One should be able to show that subsystems of statistical mechanical systems behave as if they were classical particles, taking appropriate limits from a more fundamental description of the system.

There is unfortunately rather little discussion of this position in the philosophical literature. The most extensive is a couple of short sections in (Sklar, 1993), but the points raised therein are either not relevant or are not substantive enough to engage with in detail here beyond a few quick remarks. Sklar briefly discusses the possibility of renouncing an ontological reification of statistical mechanical microstates, as I advocate, but he calls such a move a "radical proposal for revising our basic ontology" (Sklar, 1993, 363) and quickly dismisses it. Yet, our present basic ontology surely does not come from classical particle mechanics; the world is, if anything, fundamentally quantum mechanical, so no revision of our "basic ontology" is actually required.

[1] I set aside in this paper the cluster of interpretations that are characterized by their subjectivist or essentially epistemic character, especially ones based on indifference principles or ignorance. I direct the reader to substantial criticism elsewhere (Albert, 2000; Loewer, 2001; North, 2010; Meacham, 2010). (Uffink, 2011) provides a friendly and thorough review of subjectivist interpretations.

[2] To be sure, this idea is not entirely novel since similar views have been suggested before, most notably by Prigogine (See Prigogine (1996) for a popular account of his views). My argument however does not depend at all on the dynamical considerations which Prigogine argues makes a realist interpretation of phase points untenable. See (Batterman, 1991) for a philosophical evaluation of Prigogine's approach. The general view which I advocate here is discussed in the philosophical literature, as far as I am aware, only by Sklar (1993, 7.IV.2, 9.III.1) (he pejoratively calls it the "radical" or "revisionist" ontological approach), and he appears to take Prigogine alone as the source of the interpretation. Most of his substantive concerns address Prigogine's idiosyncratic views, and so they will not be discussed here.

One might wave off this objection, pointing out that the same issues in classical statistical mechanics arise in quantum statistical mechanics. There is a crucial difference between the theories however: a quantum statistical state is nothing more than a particular quantum state. Quantum statistical mechanics concerns certain, special quantum mechanical systems, whereas classical statistical mechanics, by adding a probabilistic component, is a very different theory indeed from classical mechanics (this despite the fact that physicists write quantum statistical states as if they were probabilistic weightings of pure quantum states). Thus there is no potential intra-theoretic question to be asked at all about the relation of macro- and microstates in quantum statistical mechanics as there is in classical statistical mechanics. There are no macrostates in quantum mechanics. For this reason I confine further discussion to classical statistical mechanics.

With these several preliminaries set forth, the paper will now proceed as follows. Since there is a strong default presumption against the kind of view I am advocating, I first prepare the way by arguing that a view like mine should in fact be favored by default over the traditional ones on general methodological grounds (§2). Then, I discuss the possibilities of characterizing the stochasticity of statistical mechanics, and show how various interpretations can be suggested by the formalism of the theory, including the view which I advocate (§3). I next address potential concerns with my account centering on the representational significance of phase space. In responding to these concerns, I argue that some structure of phase space, particularly phase points, should be interpreted instrumentally in statistical mechanics (§4), using as an example the phase space formulation of quantum mechanics, where quantum states are representable on phase space despite not possessing precise classical microstates. In §5 show how a commitment to a fundamental ontology of classical particles can coherently be avoided by treating statistical mechanics as a "special science" (Callender, 2011) with a structuralist ontology (Ladyman and Ross, 2007). I conclude (§6) with some brief comments to suggest potentially fruitful extensions of this work, for example to non-equilibrium statistical mechanics and quantum mechanics.

2 Interpreting Classical Statistical Mechanics

It is widely presumed that the foundations of statistical mechanics is an enterprise centrally concerned with completing the project of the theory's founders, namely of reducing thermodynamics to the mechanical motion of atoms and molecules (Sklar, 1993; Lebowitz, 1999; Callender, 1999, 2001; Frigg, 2008; Hemmo and Shenker, 2012)—"naught but molecules in motion" in Maxwell's memorable phrase. In discussions of classical statistical mechanics, whether in foundational contexts or otherwise, philosophers and physicists alike invariably adopt (in some way or another) the traditional presupposition of a classical ontology of particles that move about in space. After all, since the formal framework of classical statistical mechanics is "built on top of" that of classical particle mechanics, it seems completely natural to appropriate the particulate manner of speaking from the latter. In most popular interpretations, moreover, individual statistical mechanical systems are presumed to possess a classical mechanical state (a "microstate") that evolves deterministically and has determinate physical properties, just as such microstates do in classical particle mechanics.

The notion that microstates and their attendant ontological interpretation play an important role in statistical mechanical systems is a common denominator of the foundations of statistical mechanics, even despite the oft-acknowledged diversity of approaches taken in the field. From Boltzmannian-inspired approaches, in which individual statistical mechanical systems actually possess individual microstates, to the ensemble interpretation, in which a system's macroscopic state is understood to be a probabilistically-weighted infinite ensemble of classical mechanical microstates, nearly all ascribe some interpretive significance to microstates and the classical ontology thereof.

Statistical mechanical probabilities themselves are introduced to classical particle mechanics in essence by associating probability distributions to particular aggregations of microstates. In general I will call such probability distributions "macrostates". These macrostates, in concert with the relevant observables, completely represent the predictable content of statistical mechanics in the form of statistics (expectation values, variances, etc.). The full empirically-verifiable content of the theory, at least, resides in the macrostates of systems, not in any individual microstate (as it is in classical mechanics).

Individual microstates, by contrast, are essentially irrelevant to making predictions in the theory, as well they should, since it is generally supposed that individual systems' precise microstates are epistemically inaccessible. This is a basic assumption of statistical mechanics.[3] The most cited reason why is that quintessential statistical mechanical systems have a large number of degrees of freedom, such that it is at least *practically* impossible to adequately determine their classical microstates. If these microstates, contrary to fact, were epistemically accessible, then probabilities would simply not have to be introduced in the theory and we would be able to use classical mechanics to describe such systems (at least in principle).

Thus on the one hand, macrostates (and therefore the probabilities associated with them) tell us everything we can know about a system's empirical content (from the point of view of statistical mechanics), while having no influence at all on the behavior of individual microstates, as microstates are understood to evolve deterministically and therefore in ignorance of this probabilistic element. On the other hand, microstates, representing as they do the complete microphysical properties of a system, determine (ontologically speaking) the system's macroproperties, but they obviously cannot determine the correct probability distribution, and hence macrostate, to associate to the system, since a macrostate is a *collection* of microstates. Macroproperties are therefore over-determined in any interpretation which attributes representational significance to both macrostates and microstates. In traditional interpretations of statistical mechanics (at least those which ontologically reify microstates) this circumstance is often said to be puzzling.[4] Some philosophical work is clearly required to dissolve the tension if such interpretations are to be successful, and recognition of this challenge has indeed led to numerous efforts at establishing a philosophical account of "objective chances" to help resolve it (Loewer, 2001; Hoefer, 2007; Maudlin, 2007; Schaffer, 2007; Ismael, 2009, 2011).

Why bother to posit an ontology of classical particles, when such a posit is not only well beyond our epistemic reach (in principle even!) but, more importantly, conjures significant interpretive difficulties? I suggest that this predilection toward a particular non-fundamental basic ontology of classical particles for statistical mechanics is—at the present—inadequately justified; rather it seems to me that there should even be a prejudice against it. Here is one brief argument. First of all, I take it to be a plausible, defeasible methodological principle that "ontology recapitulates epistemology." By this slogan I mean that we should accept some (slight) degree of anti-metaphysical "positivism" in scientific metaphysics, namely by only accepting the existence of unobservable entities in a theory when we have substantive empirically-based reasons to believe that they exist (for example, the entities are in some sense "essential" in explaining some empirical phenomenon). In this vein, I note that since a non-fundamental theory may be micro-ontologically multiply-realized, it is not enough to demonstrate a handful of cases where a particular microscopic ontology fits the theory and then claim that the theory as such has that microscopic ontology in *all* cases. I claim that the interpretive stance in which statistical mechanical microstates in all cases faithfully represent actual configurations of classical particles in motion flouts this principle, in particular due to the degree of empirical inaccessibility of the supposed classical microstates and the fact

[3]Although some textbooks do not explicitly mention the term "microstate", those that do often point out this fundamental epistemic assumption: "A macroscopic physical object contains so many molecules that no one can hope to find its dynamical state by observation" (Penrose, 1970, 2); "The definition [of microstate] requires one to know the initial positions and velocities of all n particles and to follow these motions for all time. Since n is typically of the order of 10^{24}, this is of course impossible" (Ellis, 2006, 65).

[4]"If the laws are deterministic then the initial conditions of the universe together with the laws entail all facts—at least all facts expressible in the vocabulary of the theory. But if that is so there is no *further* fact for a probability statement to be about" (Loewer, 2001, 610); "The fundamental problem with understanding the probabilities in statistical mechanics to be objective is that we are meant to posit a probability distribution over a set of possible initial states, while we suppose, at the same time, that in fact only one of these initial states actually obtained" (Winsberg, 2008, 873).

that we understand quantum mechanics as ultimately fundamental. Therefore, for the sake of scientific respectability, the microstate reifier must take on the burden of explaining why the ontology of statistical mechanics diverges from the epistemology of statistical mechanics. I certainly do not suggest that such explanations cannot be provided—indeed there are many extant accounts (of variable respectability)—but I do insist that having to take on this burden is a significant interpretive cost in comparison to a view that does not create such interpretive problems from the start.

The interpretation which I advocate conceives of the foundations of statistical mechanics as concerning more than its traditional historical project. The starting point is to observe that the theory of statistical mechanics concerns statistical mechanical systems, by which I mean all those systems, like boxes of gas and finite regions of the early universe, that are well-described by the formal framework and concrete interpretive resources of the theory. In my view, the theory describes systems in terms of statistical mechanical states—probability measures—and stochastic observables, the combination of which can be used to derive accurate statistical predictions of observational outcomes. For this reason we have good reason to take seriously the interpretational significance of probability measures and observables, but not microstates. There is, moreover, no analogous mystery of what probability is doing in the theory as there is in the traditional accounts. Probability measures merely describe the actual stochastic nature of statistical mechanical systems, which fact is responsible for the statistical observable outcomes of said systems. If one wants an explanation of this stochasticity, it will have to come from another, more fundamental theory; in statistical mechanics it is a fundamental postulate.

Therefore, on the basis of what should, I think, strike philosophers of science as reasonable metaphysical methodology in the sciences, my proposed interpretation should come out somewhat ahead of the traditional ontological presuppositions. I certainly do not take this to be a decisive point against other interpretations however. There are—I repeat for emphasis—reputable solutions to the issue I raise. My aim in this section is simply to push back against the unreasonable default presumption in favor of the traditional ontological stance, a presumption which has prevented, I believe, a view like the one I am proposing to have a fair hearing. Indeed, I urge that there are good reasons to favor the basic idea of the interpretation, insofar as one accepts the plausible methodological principle that ontology should recapitulate epistemology. Whether the interpretation that follows this prescription is the better interpretation, however, depends on its precise details, to which I now turn.

3 Probability in Statistical Mechanics

A perspicuous formulation of statistical mechanics should be able to tell us what probabilities do in the theory. Statistical mechanics incorporates probability theory as a way of deriving statistical predictions of observables. Precisely how probability is implemented formally in the the theory is, however, quite flexible, as a consequence of which there exist physically-motivated alternative interpretations that identify different objects to which probability numbers are attached and different elements of the framework as random variables—in short, what is stochastic in the theory. To be sure, there is a common view on the matter, namely that probabilities attach to microstates of the system and observables are represented by random variables. My aim in this section is to demonstrate that there exist plausible alternatives to the standard view based on two modifications: (1) one where probabilities are understood to attach to observable outcomes; (2) one where microstates are understood as random variables.

Before satisfying this aim, I should say a few things about the interpretation of probability in order to forestall any confusion about the scope of my discussion. Indeed, among the most persistent questions in the philosophy of statistical mechanics is, "How are probabilities in statistical mechanics to be understood?" Especially in recent years, there have been many studies that have taken up this interpretational question (Sklar, 1993; von Plato, 1994; Guttmann, 1999; Clark, 2001; Lavis, 2001; van Lith, 2001; Emch, 2005; Uffink, 2007; Frigg, 2008; Winsberg, 2008; Meacham, 2010; Myrvold, 2016). The usual approach taken is investigating whether some one of the familiar so-called "interpretations of probability" makes sense of the application of probability in the theory. On the whole, the conclusions of these various studies have not been encouraging. Certainly some accounts of probability receive more attention than others, yet the debate continues and there remains a distinct lack of real consensus on the physical, metaphysical, and conceptual significance of statistical mechanical probabilities.

In presenting my account of statistical mechanics I adopt a restricted approach to the interpretation of probability in statistical mechanics in order to avoid this admittedly complicated morass. I also wish to avoid any distraction over the meaning of probability, so I will take probability in physics to be essentially a theoretical concept (Sklar, 1979), like "hamiltonian" or "gauge transformation". By an interpretation of probability in a specific physical theory, I will therefore mean something much narrower than usual, namely an account of what is randomly determined in the theory—in other words, an account of the theory's stochasticity. As said, what is stochastic about a theory can be usefully understood by investigating the answers to two questions:

1. To what are probabilities attached in the theory?

2. What are the random variables in the theory?

The alternative interpretation of statistical mechanics which I propose locates the stochasticity of the theory in the system's observables, i.e. the observable properties of the system are what is random about a statistical mechanical system. This interpretation can be usefully understood as stating that probabilities are attributed to observable outcomes, where the observables are treated as random variables on the space of possible observable outcomes. The precise details of these statements will emerge in the following. Already, however, one can easily see that, according to this interpretation one must reject (on pain of contradiction) the idea that the microstate of the system be understood as the instantaneous state of a large collection of particles; instead one must take seriously the idea that probability measures represent the complete physical states of individual statistical mechanical systems.

This latter claim should not be misunderstood—probability measures only represent the physical states of individual statistical mechanical systems *qua* statistical mechanical systems. I emphasize again that there certainly may be other accurate descriptions of the system in the terms of some microphysical theory. Indeed, depending on one's reductive inclinations, there perhaps must be. It is not necessarily my aim, however, to take a stand here on precisely what inter-theoretic relations there may be between statistical mechanics and more fundamental theories like quantum mechanics. I do wish to point out that this interpretive view directs attention precisely to investigating such relations however.

In any case, to the details. A general statistical mechanical system is conventionally and conveniently described kinematically by three things: a phase space Γ, a (macro-)state ρ, and a set of observables $\mathcal{A}$. The notion of a phase space is borrowed from classical particle mechanics, where it is the complete space of possible states of some system. Encoded in a classical particle state x of Γ is the spatial configuration of all individual particles as well as their states of motion, i.e. their positions as well as their momenta. In statistical mechanics the elements x of Γ are referred to as microstates, and they are often interpreted further as representing the epistemically-inaccessible but real underlying microphysical states of motion of the system. The macro-state ρ is a probability distribution $\Gamma \to [0, 1]$, and the observables $A_1, A_2, A_3, \ldots \in \mathcal{A}$ are random variables on Γ, usually represented as maps $\Gamma \to \mathbf{R}$.

A statistical mechanical system described as such naturally describes a probability space, where Γ is the sample space, the set of events $\mathcal{L}$ is conveniently taken to be the set of (Lebesgue) measurable subsets of Γ, and ρ is used to define a probability measure μ_ρ:

$$\mu_\rho(U) = \int_U \rho \, \mathrm{d}\Gamma, \tag{1}$$

where U is an element of $\mathcal{L}$, and Lebesgue integration is with respect to the natural volume element $\mathrm{d}\Gamma$ on

Γ. It is clear, when the theory is formulated this way, that probabilities ultimately attach to microstates, since ρ is understood to be a map from Γ to $[0, 1]$.

At this point I must limit the scope of my discussion somewhat for the sake of simplicity. Since statistical mechanics encompasses a fairly diverse array of approaches and formalisms (and also to avoid difficult and distracting complications), I restrict my attention to equilibrium statistical mechanics. This simplifies matters considerably, since one then does not need to worry about the complicated dynamical behavior of non-equilibrium macrostates. The case for an alternative interpretation of probability made on the basis of equilibrium statistical mechanics alone is of sufficient interest, I hope, without immediately considering statistical mechanics in all its diversity.

In general, the empirical content of statistical mechanics is given by statistics of the observables, that is, expectation values, variances, and so on (Wallace, 2015). In the case of equilibrium statistical mechanics we do not have to worry about the temporal behavior of observables or macrostates. Thus the expectation value of an observable A is given by simply treating A as a random variable on Γ associated to the probability measure μ_ρ:

$$\langle A \rangle_\rho = \int_\Gamma \rho \, A \, \mathrm{d}\Gamma. \tag{2}$$

Most approaches to equilibrium statistical mechanics can be made to fit this basic formal framework. Inspecting the previous expression, one has probabilities associated to microstates and treats observables as random variables on phase space. One integrates over the observable values at each point in phase space weighted by the probability associated to that point to determine an expectation value. But what exactly about the system is random? Is it the microstates or the observables? Surely the answer is not viscerally apparent—some intellectual work is required to explain what is happening physically.

One interpretation, common to Boltzmannian approaches to the subject, is that any statistical mechanical system possesses in addition to its macrostate ρ a classical microstate $x(t)$ (at each time t) which evolves deterministically, as in classical particle mechanics. The only thing that could be random about a statistical mechanical system which is also a deterministically-evolving classical mechanical system is its initial microstate x_0. That is to say, for any statistical mechanical system there is (or was) a random trial that determined the initial microstate x_0 of that system, which state determines its future evolution and observable properties just as in classical particle mechanics. Of course, as I argued in the previous section (§2), one must explain the overdetermination of macroscopic observables, since a statistical mechanical system's microstate is presumed to be epistemically inaccessible and therefore uncertain. The natural response to this problem is that one can only compute the statistical consequences of this epistemic uncertainty in order to derive predictions of the macroscopic observables. Thus expectation values should in fact be viewed epistemically, whereas the system's macroscopic observables are actually and fully determined (ontically) by classical mechanical observables, i.e. real-valued functions on phase space, in particular of the system's classical trajectory: $A(t) = A \circ x(t)$.

A second interpretation is possible where microstates are understood to represent the state of a system of particles. In this case microstates do not evolve deterministically as in classical particle mechanics. Instead the microstate is itself taken to be stochastic. This can be made formally more manifest by treating the microstate as a random variable on phase space, and representing the observables in the same way as observables in particle mechanics, again, namely as real-valued functions on phase space.[5] Specifically, define a random variable $X : \Gamma \to \Gamma$ on phase space which is just the identity map between measure spaces Γ. Then the expectation value of an observable can be written

$$\langle A \rangle_\rho = \int_\Gamma \rho \, (A \circ X) \, \mathrm{d}\Gamma. \tag{3}$$

[5]It seems that this view may be read into some statements in the physics literature concerning what is happening physically: "As time passes, the system continually switches from one microstate to another, with the result that, over a reasonable span of time, all one observes is a behavior 'averaged' over the variety of microstates through which the system passes" (Pathria, 1996, 30). This story is consistent with the first interpretation, but with some imagination it evokes the second interpretation as well.

There is yet another possible interpretation of the expectation value of equation (2): the observables themselves are stochastic. To make sense of this claim, however, one has to give up on the idea that statistical mechanical systems possess deterministically evolving microstates, since if the observables were truly stochastic, then the system would have contradictory observable properties (the observables realized through the stochastic process and the observables determined by the microstate). This was true of the previous interpretation as well, but in the present case I emphasize that there is additionally no need to give a realistic interpretation of the microstates. It is this view of the stochasticity of statistical mechanics which I advocate here, since it takes seriously the idea that statistical mechanical states are probability measures. It also makes no commitment about a statistical mechanical system's microphysics (qua statistical mechanical system).

My claim that there is no need to give a realistic interpretation of the microstates may strike some as puzzling, since expectation values are computed by explicitly quantifying over microstates in equation (2). It may seem like phase space is doing some important representational work, in other words. This objection does not seriously threaten the interpretation however. First, insofar as one thinks that realism is not an all or nothing affair with respect to a given formal framework—that is, insofar as one is a selective realist, instrumentalist, etc.—one is well within one's rights to interpret some aspects of a theory's formalism instrumentally. A selective realist may be completely satisfied to interpret phase space and its microstates as ontologically insignificant representational structure, while taking seriously macrostates and observables as representing real structure. The objector likely would wish to see a justification for treating phase space in such a deflationary way of course, and so I supply this in the following section. Before turning to that, though, I think it is worthwhile to demonstrate first that the burden of explaining away phase space is not necessarily forced upon one advocating the stochastic observables interpretation.

It is in fact possible to mold the above formalism somewhat so that the interpretation of observables as stochastic becomes even more manifest. This formal modification is not at all necessary in my view, but it is perhaps instructive. Just as one can treat phase space as a probability space and the microstate as a random variable on this space (as in the second interpretation), one can treat the image set of any observable as a probability space and the observable outcomes as random variables on these spaces. By this mild subterfuge one can then remove all reference to phase space and its microstates.

Given the probability space $(\Gamma, \mathcal{L}, \rho)$ it is trivial to construct these spaces. Let A be an observable of a statistical mechanical system. The measurable space associated with it is the image set $A[\Gamma]$ of A with measurable sets $\mathcal{L}_A = \{A[U]\}$, for sets $U \in \mathcal{L}$. Let ρ be the state of the system with respect to the measurable space $(\Gamma, \mathcal{L})$. The state ρ_A of the system with respect to measurable space $(A[\Gamma], \mathcal{L}_A)$ is given by $A_*\rho$, the pushforward of ρ under A.[6] One now has constructed a new probability space $(A[\Gamma], \mathcal{L}_A, \rho_A)$.

Expectation values of the system are then computed according to the following formula:

$$\langle A \rangle_{\rho_A} = \int_{A[\Gamma]} \rho_A \, A \, \mathrm{d}A. \tag{4}$$

The statistics predicted from formalizing statistical mechanics in this way, namely on probability spaces $\mathcal{A}[\Gamma]$ of observables, is of course identical to the statistics predicted from formalizing it on phase space (as a probability space). It is so by construction. Moreover this set of probability spaces manifestly suggests the stochastic observable interpretation, an interpretation which is perhaps less apparent in the standard phase space formulation. In this interpretation probabilities are clearly attached to observable outcomes, and those observable outcomes are treated as random variables. Together they formally represent the stochasticity of the observables and thereby the stochasticity of statistical mechanics manifestly.

It may, however, have occurred to the skeptical reader that the outcome probability spaces just constructed, that is those of the form $(A[\Gamma], \mathcal{L}_A, \rho_A)$, are parasitic on the phase space probability space and are therefore somehow objectionable. That is not really the case though, since one could just as well work

[6] You can think of this map as using A to pull back measurable sets of the image set of the observable to Γ and then using ρ to compute the set's probability.

the other direction, i.e. by choosing ρ_A and pulling this function back to some abstract space (or even phase space). Certainly the pullbacks are not going to be unique in general, but what is there to worry? One just has to insure that the pullbacks of the set $\rho_{\mathcal{A}} = \{\rho_A, \rho_B, \ldots\}$ agree appropriately in the target space (whatever that may be, phase space or otherwise), and then one has the resources at hand to do the full business of statistical mechanics. Presented this way, there is no obvious reason at all to demand an explanation for why statistical mechanical systems are *not* represented on phase space—after all, should one demand some such an explanation for every possible abstract unifying space to which one could pull back the observable measures? What makes phase space special in this respect, other than that it is the state space of classical particle mechanics? So far as I can see the answer is: nothing much.

4 Phase Space and Statistical Mechanics

I suspect many may nevertheless feel that there is a deep significance in the fact that statistical mechanics is generally formulated on phase space rather than some other abstract space that unifies the set of observables.[7] Some perhaps would find it auspicious enough that the alternative interpretation I offer, namely in terms of stochastic observables, appears deliberately heretical in this respect. I certainly do grant that the phase space probability space is instrumentally useful, especially since having a single state on it unifies the observable content of the system in a single, familiar space, one which can also be used for spatial constraints, etc. But that is all a physical state is supposed to be anyway: a unified representation of the observable content of the system. Once one has such a state space, it is a separate, interpretive matter to determine whether one should use phase space or some other probability space with enough representational content to coordinate the observational content of the physical system. I am happy to use phase space for these purposes, while not taking it seriously representationally.

There are, however, two further issues about phase space which I address in this section. First, I will make good on the justification for instrumentally interpreting phase space points by presenting an analogous case in quantum mechanics, where a phase space representation is also possible of a theory that is normally formulated on a different abstract space, in particular on a Hilbert space. Second, if phase space should not necessarily be taken too seriously for purposes of ontological interpretation, as I suggest, then why might phase space be nonetheless a good way to represent statistical mechanical states? My view here is that phase space has enough representational structure that it can usefully represent not only statistical mechanical states but underlying microphysical structure as well, *whatever that may be*. To explain its utility in this latter respect one would wish to have a microphysical theory that can represent its own states as probability measures on phase space—in the event, the phase space representation of quantum mechanics affords one precisely this possibility too, and therefore there is at least one possible account for phase space's utility in statistical mechanics.

Note that classical particle mechanics, the theory typically thought of as providing the underlying microphysics of classical statistical mechanics, cannot realize general macrostates in this way, for a classical mechanical state on phase space is merely a "sharply-peaked" distribution and such states cannot combine to form objective "superpositions" in the way needed to generate statistical mechanical states. Thus equilibrium statistical mechanics is naturally seen as a *more* general theory than classical mechanics, in that it can represent classical mechanical states as well as a variety of other statistical states within its scope. This turns the conventional attitude on its head: statistical mechanics, from this point of view, is the more general theory of which classical mechanics is a special case, rather than classical mechanics being the fundamental theory to which statistical mechanics is a special (and peculiar, because of the introduction of probabilities) case.

[7] "In statistical mechanics, we have, [unlike the case in quantum mechanics], built into the very mathematics from which the probability distributions are derived the underlying phase space with its pointlike representatives of exact micro-states" (Sklar, 1993, 291).

Anyway, the point is that if one wishes to understand from where statistical mechanical states come, then one should look elsewhere for the provision of appropriate phase space representations. Classical mechanics cannot provide the requisite states. As a more fundamental theory, quantum mechanics is obviously an appealing place to look. As it turns out, there is a relatively well-known, straightforward, and illuminating way to represent quantum states and operators on phase space, namely through the use of Wigner functions and Weyl transforms (Kim and Noz, 1991; Case, 2008). Let ρ be a density operator that represents a (mixed) quantum state, that is,

$$\rho = \sum_s p_s |\Psi_s\rangle\langle\Psi_s|, \tag{5}$$

where p_s is the probabilistic weighting of a particular pure quantum state $|\Psi_s\rangle$ in a mixture of pure states indexed by s. The Wigner function W_ρ is a representation, a "smearing", of such a quantum state ρ as a function on phase space. It is a map $W_\rho : \Gamma \to \mathbf{R}$ which may be defined as follows:

$$W_\rho(q,p) = \frac{1}{\pi^{N/2}} \int \mathrm{d}\Gamma\, e^{-2ipx/\hbar} \langle q - x \,|\, \rho \,|\, q + x \rangle, \tag{6}$$

where local coordinates q (position) and p (momentum) have been introduced on $2N$-dimensional phase space Γ.

Quantum mechanics on phase space can also provide the means to connect quantum mechanical microstates to statistical mechanical macrostates. If the Wigner function is a probability measure on phase space (it is not necessarily so), then it may (at least formally) serve as a "statistical mechanical state." Quantum mechanical observables may be represented on phase space similarly to quantum mechanical states (such transforms are called Weyl transforms), in particular as functions on phase space—that is, statistical mechanical observables. The expectation value of an operator A, then, is simply given by the familiar, obvious integral over phase space as in statistical mechanics:

$$\langle A \rangle_\rho = \int_\Gamma W_\rho\, A \,\mathrm{d}\Gamma. \tag{7}$$

Quantum mechanics is generally formulated in such a way that some Hilbert space is the state space of a system described by the theory, i.e. the density operator ρ is an element of a Hilbert space $\mathcal{H}$. By making use of Wigner functions and Weyl transforms Hilbert space quantum mechanics can be re-represented as a quantum theory on phase space. To be sure, merely in virtue of a phase space representation of quantum mechanics existing, one is clearly not led in this case to interpret quantum mechanical states as possessing classical microstates which correspond to phase space points. Of course one may try to do so, in which case one may then end up with something like the Bohmian interpretation of quantum theory. Nevertheless, in typical observer-independent interpretations of quantum mechanics the putative microstates, i.e. phase space points, would naturally be treated instrumentally in quantum mechanics on phase space, while the quantum state itself is taken to be real. What is true with quantum mechanics is true with statistical mechanics too—or so I suggest: phase space may be a useful representational structure for the theory for some incidental purposes, in which case it would not force any particular ontological consequence for the microphysics underlying the theory.

I give the example of quantum mechanics on phase space not to make the case that this is precisely how statistical mechanical states come about in general (although that is a quite plausible case). I mean to show rather that phase space is reasonably a vehicle for representing underlying microphysical structure in an indirect way and, therefore, that one is not forced in any way to interpret the "fundamental" ontology of the theory to be classical particle states. Moreover, there is then no auspiciousness about the appearance of phase space in the standard formalism of statistical mechanics. The fact that quantum mechanics is representable on phase space, even though quantum mechanics is not necessarily a theory of classical microstates, suffices to establish this point.[8] Furthermore, it is also not at all necessary to reformulate statistical mechanics to avoid referring to phase space (as I did in §3). Phase space is, I expect, actually of significant practical and heuristic value in establishing the sort of *inter*-theoretic relations which are needed to illuminate the foundations of statistical mechanics. It does not, however implicate a particular *intra*-theoretical ontological relation between microstates and macrostates. To suppose that it does is simply to make an interpretive mistake about representational structure.

5 Alternative Foundations

In traditional interpretations of statistical mechanics some intra-theoretical foundational relationship is usually assumed between statistical mechanics and classical mechanics (Wallace, forthcoming). One idea is to take classical particle mechanics as a fundamental theory[9]—as at least more fundamental than classical statistical mechanics—and statistical mechanics as a "higher level" theory that successfully describes a particular set of phenomena—thermodynamic phenomena—despite epistemic limitations in accessibility to the "lower level" details. It is also sometimes supposed, however, that statistical mechanics is a fundamental theory of sorts, and that the probabilities are somehow objective features of an otherwise classical mechanical world of particles moving in space (Albert, 2000; Loewer, 2001).

The assumption of a necessary relationship between classical and statistical mechanics in these ways should strike us as somewhat odd—again, the world is quantum mechanical, so far as we know, not classical mechanical, and therefore there is a multitude of possible inter-theoretic relations between these three theories to consider besides two simplistic, direct reductions.[10] Assuming that the ontological foundation of classical statistical mechanics is the same as found in classical particle mechanics not only faces the epistemological challenge raised in §2, but it is entirely unnecessary at the present, since there are these days other potential fundaments for the former besides the latter. So why assume that *all* statistical mechanical states reduce to classical particle states?

A key virtue of adopting the interpretation of statistical mechanics which I am advocating is that it avoids any unwarranted commitment to a particular underlying ontology and foundation for the theory. Sklar finds this "revisionist" approach to the ontology of statistical mechanics objectionable precisely because it fails to complete the historical project of statistical mechanics. As he says, "denying the existence of the exact micro-states...doesn't seem to help us at all in resolving the most crucially puzzling questions such as those centered around the origin of time-asymmetry" (Sklar, 1993, 366). On the contrary, I believe denying the necessity of a particle ontology for statistical mechanics can help solve those problems, namely by re-directing foundational efforts from the fruitless search for explanations mired in overly classical thinking to other avenues of explanation based on our best philosophical understanding of current physics (much of which is quantum mechanical). Avoiding unnecessary ontological commitments does not force one to be agnostic about foundational matters; rather the interpretation encourages the fruitful and entirely sensible reorientation of foundational investigations into statistical mechanics from the *intra-theoretic* to the *inter-theoretic*.

No doubt severing the traditional firm connection between statistical mechanics and classical mechanics leaves one with some questions. How can a *probability measure* represent the state of a system? How

[8] "In quantum mechanics, the mathematical apparatus posits no further underlying 'point' phase space state of the system beyond the probability distribution. Indeed, classical phase space is rejected altogether and is replaced by a phase space in which the mathematical representatives of the probability distributions (or, rather, the probability amplitudes from which the probability distributions are constructed by multiplying one of them by its complex conjugate—Hilbert space vectors) are the 'atoms' " (Sklar, 1993, 291).

[9] 'Mechanics is a completely general theory, that is it ought to give a complete description of any physical situation to which it applies' (Clark, 2001, 271).

[10] 'Why should we consider quantum issues when working in the foundations of statistical physics? The simple (too simple) answer is that classical physics is false. If our purpose, in doing foundational work, is to understand the actual world, it is necessary to use a theory which validly describes that world' (Wallace, 2001, 1).

can this alternative interpretation coherently avoid a commitment to some particular foundational viewpoint? To some extent these questions have been answered in the preceding discourse, but for lingering doubts I offer a couple of connections to recent philosophical work that I hope illuminate this account further.

One appealing way to understand statistical mechanics is as a "special science" (Callender, 1997; Callender and Cohen, 2010; Callender, 2011). For Callender statistical mechanics is a theory like those found in biology or economics:

> "It's a new special science, one that grounds and unifies a lot of macroscopic behavior. It too is restricted to certain kinds of systems, in this case macroscopic systems whose energy and entropy are approximately extensive. (Surely a better characterization can be given, but this will do for now.) The claim is that the [statistical mechanical] probabilities only kick in when we have systems meeting such a description. Once they develop, one uses [statistical postulates] to great effect. But one does not use it to describe and explain the frequency with which such systems develop in the first place. The science of statistical mechanics is about systems with certain features. Yet it's no part of the science to say anything about the frequencies of these systems themselves" (Callender, 2011, 110).

This is just to say that it is a job for another theory with greater explanatory resources to account for the particular actual states of the theory that arise in the world. Indeed, it may be a job for other *theories*, since it is certainly open that the probability measures which serve as statistical mechanical states may connect to microscopic structures in a variety of ways. In any case, by treating statistical mechanics as a special science, this interpretation of statistical mechanics avoids an unnecessary commitment to a particular intrinsic foundation.

Structuralism encourages the further view that it is unnecessary to posit any additional ontology for the theory over and above the structure given by the theory itself (Ladyman and Ross, 2007). Structuralism, in its most benign form, should certainly not be seen as an effort to "get rid of" ontology altogether but rather as an effort to proof our metaphysical thinking of ontological prejudices—for example that the world is necessarily made up of individual classical particles or that all statistical mechanical systems are composed of elements which can be described in classical particle mechanics.[11]

Such prejudices can impede progress in understanding the foundations of physics and may even impede physics beyond foundational matters. As Cao (2003, 5-6) observes, "an ontological commitment made in a scientific discipline also dictates its theoretical structure and the direction of its evolution." An ontological commitment to classical particles dictated the theoretical structure of statistical mechanics, and has since dictated the direction of its evolution within both physics (including, to some extent, quantum statistical mechanics) and philosophy. We now know much more about microphysics than the founders of the theory, and therefore a foundational account that presupposes a particle ontology should, it seems to me, be generally viewed as the depreciated metaphysical fiction of a former time (which admittedly was and to some extent remains of significant heuristic value however!). The alternative account proposed here provides one coherent starting point to move forward beyond it in philosophical and foundational endeavors.

In short, it is in a structuralist way that I suggest one understand probability measures as statistical mechanical states. At the level of description of physical systems captured by statistical mechanics there is, I claim, an unavoidable stochastic element which is naturally represented mathematically by probability measures. This is what it means, in a sense, to take probability measures seriously as statistical mechanical states, with the probabilities representing the stochasticity of the system's macroscopic observables. Just as one can understand quantum states as probability measures (via Gleason's theorem),

[11] "We are not 'anti-ontology' in the sense of urging a move away from electrons, elementary particles etc. and towards 'observable structures' or the S-matrix or whatever; rather, we urge the reconceptualization of electrons, elementary particles and so forth in structural instead of individualistic terms" (French and Ladyman, 2003, 37).

one can also understand statistical states as probability measures as well. The upshot of such interpretations is that physical states represented as such indicate an interesting stochasticity in the world, at least at the level of description captured by those theories. That, I urge, is a profound metaphysical discovery that should provoke further thought and investigation, rather than reversionist attempts at preserving ontological bygones.

6 Concluding Remarks

The prejudice in foundational work on statistical mechanics towards interpreting microstates as representing the states of systems of particles has unfortunately suppressed consideration of alternative interpretations of the theory and its probabilities. The only extant criticisms of this account have been overstated. Although the limited discussion of this particular alternative has so far centered around Prigogine's work, the general view which I advocate is independent of his school's specific proposals and represents a novel philosophical interpretation of statistical mechanics, one which I have argued is plausible on the philosophical and theoretical grounds promoted here. No doubt much work remains to flesh out the view more fully, but I hope to have established that it can and should be taken seriously as a viable alternative to the traditional attitudes that dominate in the foundations of statistical mechanics.

To review briefly: I have argued that to metaphysically reify the microstates of phase space is to make suppositions about an underlying ontology, suppositions that have no grounds in the empirical or even the theoretical facts of statistical mechanics. Moreover, I claim that there is no need to be alarmed at the deferral of foundational questions to other theories or the rejection of a particular fundamental ontology for statistical mechanics. If one understands statistical mechanics structurally and as a sort of special science, then there is no need for it to justify its posits or to seek answers to difficult foundational questions *within its own framework*. There is surely some story as to why statistical mechanical systems behave as they do, have the states that they do, and are well represented in a phase space picture. I have insisted that this story however plausibly involves more fundamental theories than classical particle mechanics, and hence more foundational work to establish the relevant inter-theoretic relations.

Adopting this interpretation of statistical mechanical states opens up many exciting new possibilities for foundational and philosophical work. I will mention only a couple; the reader can no doubt envision more. In the first place the possibility of interpreting statistical mechanics in this way makes for a stronger connection to the interpretation of quantum mechanics and the measurement problem, a connection at which I have only briefly hinted in this paper. There are interesting analogies (and dis-analogies) between popular interpretations of quantum mechanics and the ways of interpreting the stochasticity of statistical mechanics and its statistical predictions, some of which are more or less known, others which remain to be explored. Secondly, taking probability measures seriously as statistical mechanical states also makes it very clear that the Hamiltonian dynamics of classical particle mechanics is inappropriate for describing the dynamics of non-equilibrium statistical mechanics. Since the action of a Hamiltonian defined on phase space on probability measures defined on phase space is non-dissipative, no statistical mechanical state will ever approach equilibrium. Although particle-based intuitions lead to empirically successful non-equilibrium equations of motion like the Fokker-Planck equation (Kadanoff, 2000, Ch. 6), the interpretation I advocate suggests viewing these intuitions as merely heuristic. This interpretation may therefore suggest novel ways of understanding the utility of such equations in non-equilibrium statistical mechanics.[12]

[12]Such philosophical work would apparently be of use to physicists as well: "These equations are just convenient to use; we know how to use them but do not fully understand why they can be so useful" (Ma, 1986, 367).

References

Albert, David. *Time and Chance*. Cambridge, MA: Harvard University Press, 2000.

Batterman, Robert. "Randomness and Probability in Dynamical Theories: On the Proposals of the Prigogine School." *Philosophy of Science* 58: (1991) 241–263.

Callender, Craig. "What is 'The Problem of the Direction of Time'?" *Philosophy of Science* 64 (Proceedings): (1997) S223–234.

———. "Reducing Thermodynamics to Statistical Mechanics: The Case of Entropy." *The Journal of Philosophy* 96: (1999) 348–373.

———. "Taking Thermodynamics Too Seriously." *Studies in History and Philosophy of Modern Physics* 32, 4: (2001) 539–553.

———. "The Past Histories of Molecules." In *Probabilities in Physics*, edited by Claus Beisbart, and Stephan Hartmann, Oxford: Oxford University Press, 2011, 83–113.

Callender, Craig, and Jonathan Cohen. "Special Sciences, Conspiracy and the Better Best System Account of Lawhood." *Erkenntnis* 73: (2010) 427–447.

Cao, Tian Yu. "Structural Realism and the Interpretation of Quantum Field Theory." *Synthese* 136: (2003) 3–24.

Case, William. "Wigner functions and Weyl transforms for pedestrians." *American Journal of Physics* 76: (2008) 937–946.

Clark, Peter. "Statistical Mechanics and the Propensity Interpretation of Probability." In *Chance in Physics*, edited by Jean Bricmont, Detleff Dürr, Maria Galavotti, Giancarlo Ghirardi, Francesco Petruccione, and Nino Zanghi, Berlin: Springer Verlag, 2001, 271–281.

Ellis, Richard. *Entropy, Large Deviations, and Statistical Mechanics*. Berlin: Springer-Verlag, 2006.

Emch, Gérard. "Probabilistic issues in statistical mechanics." *Studies in History and Philosophy of Modern Physics* 36: (2005) 303–322.

French, Steven, and James Ladyman. "Remodelling Structural Realism: Quantum Physics and the Metaphysics of Structure." *Synthese* 136: (2003) 31–56.

Frigg, Roman. "A Field Guide to Recent Work on the Foundations of Statistical Mechanics." In *The Ashgate Companion to Contemporary Philosophy of Physics*, edited by Dean Rickles, London: Ashgate, 2008, 99–196.

Guttmann, Yair. *The Concept of Probability in Statistical Physics*. Cambridge: Cambridge University Press, 1999.

Hemmo, Meir, and Orly Shenker. *The Road to Maxwell's Demon*. Cambridge: Cambridge University Press, 2012.

Hoefer, Carl. "The Third Way on Objective Probability: A Sceptic's Guide to Objective Chance." *Mind* 116: (2007) 549–596.

Ismael, Jenann. "Probability in Deterministic Physics." *The Journal of Philosophy* 106: (2009) 89–108.

———. "A Modest Proposal about Chance." *The Journal of Philosophy* 108: (2011) 416–442.

Kadanoff, Leo. *Statistical Physics*. Singapore: World Scientific, 2000.

Kim, Young, and Marilyn Noz. *Phase Space Picture of Quantum Mechanics*. Singapore: World Scientific, 1991.

Ladyman, James, and Don Ross. *Every Thing Must Go*. Oxford: Oxford University Press, 2007.

Lavis, David. "The Concept of Probability in Statistical Mechanics." In *Frontiers of Fundamental Physics 4*, edited by Burra Sidharth, and Mikhail Altaisky, Berlin: Springer Verlag, 2001, 293–308.

Lebowitz, Joel. "Statistical Mechanics: A Selective Review of Two Central Issues." *Reviews of Modern Physics* 71: (1999) S346–S357.

Loewer, Barry. "Determinism and Chance." *Studies in History and Philosophy of Modern Physics* 32: (2001) 609–620.

Ma, Shang-Keng. *Statistical Mechanics*. Singapore: World Scientific, 1986.

Maudlin, Tim. "What could be objective about probabilities?" *Studies in History and Philosophy of Modern Physics* 38: (2007) 275–291.

Meacham, Christopher. "Contemporary Approaches to Statistical Mechanical Probabilities: A Critical Commentary—Part I: The Indifference Approach." *Philosophy Compass* 5, 12: (2010) 1116–1126.

Myrvold, Wayne. "Probabilities in Statistical Mechanics." In *Oxford Handbook of Probability and Philosophy*, edited by Christopher Hitchcock, and Alan Háyek, Oxford: Oxford University Press, 2016, chapter 27.

North, Jill. "An empirical approach to symmetry and probability." *Studies in History and Philosophy of Modern Physics* 41: (2010) 27–40.

Pathria, Raj. *Statistical Mechanics*. Oxford: Butterworth-Heinemann, 1996.

Penrose, Oliver. *Foundations of Statistical Mechanics*. Oxford: Pergamon Press, 1970.

Prigogine, Ilya. *The End of Certainty*. New York: The Free Press, 1996.

Schaffer, Jonathan. "Deterministic chance?" *British Journal for the Philosophy of Science* 58: (2007) 113–140.

Sklar, Lawrence. "Probability as a Theoretical Concept." *Synthese* 40: (1979) 409–414.

———. *Physics and Chance*. Cambridge: Cambridge University Press, 1993.

Uffink, Jos. "Compendium of the Foundations of Classical Statistical Physics." In *Philosophy of Physics*, edited by Jeremy Buttefield, and John Earman, Amsterdam: Elsevier, 2007, volume B of *Handbook of the Philosophy of Science*, 923–1074.

———. "Subjective Probability and Statistical Physics." In *Probabilities in Physics*, edited by Claus Beisbart, and Stephan Hartmann, Oxford: Oxford University Press, 2011, chapter 2, 25–49.

van Lith, Janneke. "Ergodic Theory, Interpretations of Probability and the Foundations of Statistical Mechanics." *Studies in History and Philosophy of Modern Physics* 32: (2001) 581–594.

von Plato, Jan. *Creating Modern Probability*. Cambridge: Cambridge University Press, 1994.

Wallace, David. "Implications of quantum theory in the foundations of statistical mechanics.", 2001. http://philsci-archive.pitt.edu/id/eprint/410.

———. "The quantitative content of statistical mechanics." *Studies in History and Philosophy of Modern Physics* 52: (2015) 285–293.

———. "Inferential vs. Dynamical Conceptions of Physics." In *What is Quantum Information?*, edited by Olympia Lombardi, Cambridge: Cambridge University Press, forthcoming .

Winsberg, Eric. "Laws and chances in statistical mechanics." *Studies in History and Philosophy of Modern Physics* 39: (2008) 872–888.