

An Alternative Interpretation of Statistical Mechanics

C. D. McCoy^{*†}

1 May 2018

Abstract

Abstract: In this paper I propose an interpretation of classical statistical mechanics that centers on taking seriously the idea that probability measures represent complete states of statistical mechanical systems. 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 are representationally significant in the theory is therefore dispensable, a consequence which suggests interesting possibilities for developing non-equilibrium statistical mechanics and investigating inter-theoretic answers to the foundational questions of statistical mechanics.

1 Introduction

A common lament 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 “we have no choice but to dwell on its history” (Uffink, 2007, 923) when investigating these foundations.

While that history has indeed witnessed many approaches to foundational questions, a recurring tension emerges in a considerable number of the debates, thematically emphasized throughout Sklar’s panoramic survey of the field, *Physics and Chance*. In statistical mechanics there seems to be a dilemma between preserving some thermodynamic regularity or fact exactly, and insuring that the theory describes individual microscopic systems (Sklar, 1993; Callender, 2001). Choosing the first horn has usually led to the adoption of an ‘ensemble’ concept of statistical mechanics, which evokes “the imaginative picture of an innumerable 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 second usually leads to giving up the exact laws in favor of a statistical regularities, in order to preserve the idea that each individual system has an exact microstate.¹

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

[†]**Acknowledgments:** Thanks to Harvey Brown, Craig Callender, Richard Dawid, John Dougherty, Gijs Leegwater, Fred Muller, Patricia Palacios, Anney Thresher, David Wallace, Stefan Wintein, and Christian Wüthrich for helpful comments and conversations on this paper. 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. Finally, I gratefully acknowledge the very helpful comments of three referees.

¹I 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,

As many commentators have noted, the ensemble interpretation of statistical mechanics makes it unclear how statistical mechanics can explain the behavior of individual systems. This has emboldened many philosophers and some physicists to favor the latter, so-called ‘Boltzmannian’, approach. Yet there is an alternative, which has received scant attention from philosophers, that parts ways with the two horns of Sklar’s dilemma. This interpretation takes seriously the idea that statistical mechanical states are correctly and completely described by probability measures. While this is to some extent already assumed in both the ensemble and Boltzmannian interpretations, something distinct is meant, namely, that statistical mechanical systems have individual states (as in the Boltzmannian interpretation) but that these states are fully represented in the theory by probability measures (as in the ensemble interpretation). It is then but a small step from this conception of state to the idea that the real, physical stochasticity of statistical mechanical systems is located in the behavior of macroscopic observables (as I will show).

Furthermore, it follows directly from the stochasticity of the observables that one cannot consistently conceive of statistical mechanical systems in the Boltzmannian way on this interpretation, that is, as possessing deterministically-evolving, microscopic, mechanical states in addition to the macroscopic state represented by a probability measure. I emphasize that this is not to say that this ‘stochastic observables’ interpretation is an instrumentalist interpretation of the theory. Plausibly there is always a microscopic state of any statistical mechanical system, one that is properly explicated in a more fundamental theory (like quantum mechanics). Indeed, my proposal of this interpretation is importantly motivated by the idea that the foundations of statistical mechanics is best explored *inter-theoretically* rather than *intra-theoretically* (as it is in traditional interpretations), for example, by investigating the classical limit of quantum systems.²

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. I stress that this interpretation does not preclude the possibility of some statistical mechanical systems appropriately reducing to classical systems of particles (in some limit). In certain circumstances reasonable cases can be and have been made. My charge is that assuming that they do in all cases is unwarranted, particularly given what we presently know about microscopic physics.

To be sure, the basic thought behind my proposal is not entirely novel. Similar views have been suggested before, most notably by Prigogine.³ The core idea of Prigogine’s program is to model statistical mechanical systems with probability distributions, as I wish to do, but he additionally aims to characterize their dynamics directly in terms of such distributions. While there is no doubt that Prigogine and his followers have made many technical contributions to statistical mechanics, it does seem that the interpretation of their work remains less than fully transparent (Bishop, 2004). Although I see affinities between some of their statements and my own view, in proposing this interpretation here I do not necessarily share their motivations, such as, for example, trying to ‘explain’ the arrow of time dynamically or incorporating chaos and stability considerations into the theory. My motivations are squarely centered on philosophical concerns, particularly having to do with the interpretation of probabilities in statistical mechanics and what ramifications alternative interpretations might have for developments of the theory.

I note also that a similar view is suggested by Krylov (1979), whose motivations are perhaps closer to mine in that he disputes the a priori adoption of the classical state as fundamental in statistical mechanics. However, I do not agree with his arguments that an interpretation of statistical mechanics with such states is impossible as a matter of principle. Unfortunately Krylov died before he could develop his positive proposal, so it is not clear how much the proposal suggested here has in common with how Krylov’s

2000; Loewer, 2001; North, 2010; Meacham, 2010); see also (Uffink, 2011) for a friendly and thorough review of subjectivist interpretations.

²This general point of view has been urged especially by Wallace (2001, 2015).

³See (Prigogine, 1996) for a popular account of his views and (Prigogine, 1962) for an early statement of the core ideas of his program. Useful discussion of the work of Prigogine’s group can be found in (Batterman, 1991; Sklar, 1993; Bricmont, 1996; Bishop, 2004).

approach might be developed.⁴

Apart from the discussion of the relatively well-known specific proposals of Prigogine and Krylov, 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). Sklar pejoratively calls it the ‘radical’ or ‘revisionist’ ontological approach. He appears to take Prigogine as the main source of the interpretation and most of his substantive concerns address Prigogine’s idiosyncratic views. Where his concerns touch on the general proposal, specifically on the renouncing of an ontological reification of statistical mechanical microstates, he mostly gives irrelevant or question begging arguments. For example, he claims that “denying the existence of the exact pointlike micro-state is neither necessary nor sufficient to solve our ultimate theoretical problems” (Sklar, 1993, 366). So it is—but who would claim so much? By the same token, neither is reifying exact pointlike microstates necessary or sufficient to do the same. Sklar also repeats a point made by Batterman, that “nothing short of [a] bona fide no-hidden-variables proof can genuinely compel one to give up the concept of an exact state” (Batterman, 1991, 260). To this point I say: one is very welcome to set the standard of one’s dogmatism wherever one likes, but insisting on it is certainly no argument against an alternative. Finally, Sklar calls the denial of such states a “radical proposal for revising our basic ontology” (Sklar, 1993, 363). 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 required at all.⁵

One might wave off this last observation, 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, whereas a classical statistical state is something quite different from a classical mechanical state. It is true that physicists write quantum statistical states as if they were probabilistic weightings of pure quantum states. Nevertheless, it is crucial to recognize that a probabilistic weighting of pure quantum states is itself just a kind of quantum state; a probabilistic weighting of classical mechanical states is, by contrast, plainly not a classical mechanical state. For precisely this reason there can be no fundamental intra-theoretic question about the relation of macro- and microstates *within* quantum statistical mechanics, as there can be in classical statistical mechanics. Accordingly, I confine further discussion to classical statistical mechanics and set aside quantum statistical mechanics, the interpretation of which is a matter of interpreting quantum mechanics in general.

The main body of the paper proceeds 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, showing 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. I conclude with some suggestions of potentially fruitful extensions of these ideas, for example, to non-equilibrium statistical mechanics and quantum mechanics.

⁴See (Batterman, 1990) and (Sklar, 1993, 7.III.5) for philosophical assessments of Krylov’s work.

⁵Sklar does raise an objection that deserves attention. He questions how a proposal like mine can deal with the results of the spin echo experiment. While I do not believe examples like this threaten my proposal when the results are suitably interpreted, it is far too specific of an example to address in a programmatic paper like this one, so its treatment (and the treatment of similar troublesome cases like Brownian motion) is best left for another paper.

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, according to which individual statistical mechanical systems possess individual microstates, to the ensemble interpretation, according to which a system's macroscopic state is a probabilistically-weighted infinite ensemble of classical mechanical microstates, nearly all ascribe some interpretive significance to microstates and the presumed classical ontology thereof.

Statistical mechanical probabilities themselves are traditionally introduced to classical particle mechanics by associating probability distributions to particular aggregations of microstates. In general I will call such probability distributions *macrostates* because of the particular role they play in the theory: the state dynamics operates on them, and, in concert with the relevant observables, they completely represent the predictable content of statistical mechanics in the form of statistics (expectation values, variances, etc.).

To forestall confusion which this characterization might invite, a few more words on what is meant here by a state (in the context of a probabilistic theory) is necessary. In classical mechanics the state of the system plays the roles just mentioned: the state dynamics operates on it, and it represents the predictable empirical content of the theory. But it also plays another role: it specifies that actual properties that the system possesses at that time. However, these roles can come apart (to some extent) in a probabilistic theory. The former roles are more significant theoretically for picking out the state of a system, which is why I choose to call probability distributions macrostates in statistical mechanics. The latter role is important only for specifying the empirical history of the system, that is, the temporal sequence of that system's realized, observable properties. One might well call this collection of properties a state too, but it does not play the other roles in statistical mechanics, except derivatively. Furthermore, the determination of a system's macroproperties cannot, in general, specify its ‘statistical state’. For this reason such a statistical state (macrostate, in my terminology) cannot be thought of as a compact description of a system's actual properties (it might instead be thought of as a compact description of its observables’ ‘potentialities’).

Individual microstates, in contrast to macrostates, 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.⁶ The most cited reason for making this assumption 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

⁶Although 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).

microstates. If these microstates 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).

So, on the one hand, macrostates (and hence the probabilities associated with them) tell us everything we can theoretically know about a system's predictable empirical content (from the point of view of statistical mechanics), while having no consequences at all for the behavior of the postulated individual microstates (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 according to the prevailing view, determine (ontologically speaking) the system's macroproperties (which are presumed to reduce to microproperties), but they cannot determine the correct probability distribution to assign to a system, as such a probability distribution assigns probabilities to a collection of (almost entirely non-actual) microstates. In traditional interpretations of statistical mechanics (at least those which ontologically reify microstates) this duality in state concept is often said to be puzzling.⁷ Some clarification is clearly required to dissolve the conceptual tension between macrostates and microstates if such interpretations are to be successful. Indeed, recognition of this challenge has led, among other things, to numerous efforts at establishing a philosophical account of 'deterministic chance' to help resolve it (Bradley, 2017).⁸

Given this basic conundrum, it is worth pausing and asking ourselves: Why should we bother with positing an ontology of classical particles at all? Why do so when such a posit is not only well beyond our epistemic reach (in principle even!) but, more importantly, creates significant interpretive difficulties? I do not think there are ready answers to these questions, despite the widespread predilection towards this particular basic ontology in statistical mechanics. As far as I can see, its assumption is poorly justified. Indeed, it seems to me that there should even be a (defeasible) prejudice against it.

Here is one argument to encourage the intuition that something is amiss in the traditional approach. I take it to be a plausible methodological principle, one which certainly admits exceptions, that 'ontology recapitulate epistemology'. This slogan is meant to evince some (slight) degree of anti-metaphysical 'positivism' in scientific metaphysics, by mandating acceptance of unobservable entities only when we have substantive empirically-based reasons to believe that they exist. Perhaps the entities are in some sense 'essential' in explaining an empirical phenomenon, for example. 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 therefore has that microscopic ontology in all cases.

The interpretive stance which holds that statistical mechanical microstates in all cases faithfully represent actual configurations of classical particles flagrantly flouts this principle, particularly given the degree of empirical inaccessibility of the supposed classical microstates and also given the fact that we regard quantum mechanics as ultimately fundamental. 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. One has an easy answer if one takes probabilities to be merely epistemic, but for those that take probabilities to be ontic the task is not an easy one. I certainly do not wish to 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 concerned with something other than its traditional historical project. The starting point is to observe that the theory

⁷For example, "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).

⁸See (McCoy, 2018a) for a critique of many of these accounts.

of statistical mechanics is about statistical mechanical systems: 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 macroscopic outcomes. For this reason we have good reason to take seriously the interpretational significance of probability measures and observables—but not of microstates. Furthermore, there is no analogous mystery of what probability is doing in the theory as there is in the traditional ontic accounts. Probability measures merely represent the actual or apparent stochastic nature of statistical mechanical systems, which fact is responsible for the statistical observable outcomes of such systems. If one wants an explanation of the origin of this stochasticity (if such is even possible), it will have to come from another, more fundamental theory; in statistical mechanics it is a fundamental postulate.

Thus, on the basis of what should, I think, strike philosophers of science as reasonable metaphysical policy 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. There are, as said, reputable responses to the tension caused by reifying microstates. My project in this section is simply to push back against an unreasonable default presumption in favor of the traditional ontological stance, a presumption which has prevented a view like the one I am proposing to have a fair hearing. I have urged that there is a reasonable motivation for the basic idea of the interpretation, insofar as one accepts the plausible methodological principle that ontology should recapitulate epistemology. Of course, if inclination is to turn into acceptance, it must rest on the interpretation's precise details, to which I now turn.

3 Probability in Statistical Mechanics

A perspicuous interpretation of a probabilistic theory should 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 formally implemented in the theory is quite flexible, however. There exist alternatives which identify the probabilistic structure of the theory differently, by identifying different objects to which probability numbers are attached and different elements of the framework as random variables. To be sure, there is a common view on the matter: probabilities attach to microstates of the system and observables are random variables. This common view informs the interpretation of the theory. My aim in this section is to demonstrate that there exist perspicuous alternatives to the standard view based on two modifications: (1) one where probabilities are understood to attach to observable outcomes, and (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 avoid any confusion about the scope of my discussion. “How are probabilities in statistical mechanics to be understood?” is, after all, surely among the most persistent questions in the philosophy of statistical mechanics. Especially in recent years, there have been many studies that have taken up this interpretational question.⁹ 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 the theory's probabilities.

Therefore I adopt a restricted approach to the interpretation of probability in order to avoid this morass. I also wish to avoid any distraction over the ‘meaning’ or ‘nature’ of probability, so for present

⁹See, for example, (Sklar, 1993; von Plato, 1994; Guttman, 1999; Clark, 2001; Lavis, 2001; van Lith, 2001; Emch, 2005; Uffink, 2007; Frigg, 2008; Winsberg, 2008; Meacham, 2010; Myrvold, 2016).

purposes 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 intend something much narrower than usual, namely, an account of what is randomly determined. I refer to such an account as an account of the theory’s stochasticity. Such an account should give answers to (at least) the following 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: that is, 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 significance of these statements will emerge below. Already one can discern that according to this interpretation one must reject (on pain of contradiction) the idea that the microstate of the system is the instantaneous state of a large collection of deterministically evolving particles. Accordingly, one must take seriously the idea that probability measures represent the complete physical states of individual statistical mechanical systems.

The scope of 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 may be other accurate descriptions of the system in the terms of a more fundamental microphysical theory. Indeed, depending on one’s reductive inclinations, there perhaps must be. It is not my brief 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, though, that this interpretive view directs attention precisely to investigating such relations.

At this point I choose to limit the scope of the discussion to equilibrium statistical mechanics for the sake of simplicity and perspicuity. As said, statistical mechanics encompasses a fairly diverse array of approaches and formalisms, each of which raises difficult and distracting complications, particularly in the non-equilibrium case. Rather than attempt to navigate this labyrinth in all of its disorienting complexity, I hope that the case for an alternative interpretation of probability made on the basis of equilibrium statistical mechanics alone will suffice. The astute reader, I trust, will be able to see, if even only dimly, how the idea sketched in the following might be generalized beyond the strictures adopted here.

To flesh out the details of the various implementations of probability and their interpretations, it is useful to adopt a simple, general formalism for equilibrium statistical mechanics. The state and empirical content of a statistical mechanical system can be described by three things: a phase space Γ , a (macro)state ρ , and a set of observables \mathcal{A} . Most approaches to equilibrium statistical mechanics can be made to fit this basic formal framework. The notion of a phase space is borrowed from classical particle mechanics, where it is the complete state space of some classical mechanical 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, that is, their positions as well as their momenta. As already noted, in statistical mechanics the elements x of Γ are also referred to as microstates, and they are usually interpreted further as representing the epistemically-inaccessible but real underlying microstates of the system. The macrostate ρ is a probability distribution $\Gamma \rightarrow [0, 1]$, and the observables A_1, A_2, A_3, \dots in \mathcal{A} are random variables on Γ , usually represented as maps $\Gamma \rightarrow \mathbf{R}$. These observables represent macroscopic properties of the system, for example its temperature and energy.

A statistical mechanical system described in this way is easily made into 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 via

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

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

Let us now focus on the representation of the empirical content of statistical mechanics. This content 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. So, the expectation value $\langle A \rangle$ of an observable A is given simply by treating A as a random variable on Γ associated to the probability measure μ_ρ :

$$\langle A \rangle_\rho = \int_\Gamma \rho A \, d\Gamma. \quad (2)$$

Inspecting this expression, one should observe that probabilities are associated to microstates (the domain of ρ is Γ) and observables are treated as random variables on phase space ($A : \Gamma \rightarrow \mathbf{R}$). To compute the expectation value one integrates over the A 's values at each point in phase space weighted by the probability associated to that point by ρ . But what exactly about the system is random? Is it something about the microstates or something about the observables? Surely the answer to this question is not given merely by inspection. Some interpretive 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 and determines all of the system's properties, 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 with properties determined by its microstate is its initial microstate x_0 . In other, more evocative words, for any statistical mechanical system, there is (or was) a 'random trial' that determined the initial microstate x_0 of that system. This state then determines the system's entire evolution and all its observable properties for all time, just as in classical particle mechanics.¹⁰ Hence, on this interpretation probabilities are in fact properly attached to *initial* microstates and are the real random variables of the theory, as observables are fully determined by the actual microstate and are not (directly) a matter of chance.

A second interpretation is possible, another one where microstates represent the state of a system of particles. In this version, however, microstates do not evolve deterministically. Instead, the microstate is itself taken to be the locus of stochasticity. As in the Boltzmannian approach, one takes the system's properties to be completely determined by the microstate, exactly like observables in particle mechanics. This idea can be made more manifest by explicitly representing the microstate as a random variable on phase space. Formally, define a random variable $X : \Gamma \rightarrow \Gamma$ on phase space that is just the identity map. Then the expectation value of an observable can be written

$$\langle A \rangle_\rho = \int_\Gamma \rho (A \circ X) \, d\Gamma. \quad (3)$$

In this case probabilities attach to microstates (now at all times, not just initial times) and are, again, properly the real random variables of the theory.¹¹

The third possible interpretation of the expectation value of equation (2) is the one I advocate: the observables themselves are the locus of stochasticity. According to this interpretation, probabilities are properly attached to observable outcomes, not microstates (as in the other two interpretations). Indeed, as I have remarked, to make sense of this interpretation one has to give up on the idea that statistical

¹⁰I do not insist, of course, that this initial random trial entails a realist commitment. One may certainly interpret initial chances in a Humean way, for example, as in (Albert, 2000; Loewer, 2001).

¹¹This view may perhaps be read into some statements in the physics literature. For example: "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). While this story is, strictly speaking, consistent with the first interpretation, with some imagination it evokes the second interpretation as well.

mechanical systems possess deterministically evolving classical microstates. This is because the system would have contradictory observable properties if it did possess such a microstate: the observables realized through the stochastic process and the observables determined by the microstate.

Cue the hue and cry: “Radicalism!” “Revisionism!” “How can *anyone* seriously maintain an interpretation of statistical mechanics that eschews microstates,” especially given the “variety and consilience of evidence” (as a referee put it) for discrete, microscopic entities well-described by classical particle mechanics? “Does he somehow aim to repudiate the atomic hypothesis with hand-waving talk of ‘observables’ and ‘stochasticity’?”

My goodness, of course not! Yet, hyperbole aside, I do admit that there are serious worries along the lines of these complaints which deserve some response in brief.¹²

Firstly, I believe people grossly overestimate the import of the alleged body of evidence for discrete, microscopic entities. That such entities exist is a *metaphysical* conclusion, for which empirical evidence can hardly be decisive. It is true that when we are able to sufficiently isolate some small part of a system we often find that we can describe it as a system of a small number of bodies. Yet it is a significant (and unwarranted) inductive leap to suppose that that same description applies to the whole undifferentiated system. Indeed, taking this leap lands one directly into the interpretive hot water discussed above. Of course one might retreat from the heat to a subjectivist interpretation of probability in order to uniformly maintain a microscopic ontology of particles. For those who wish to hang on to the physicality of statistical mechanical probabilities, however, it is difficult to see how such probabilities can matter for ‘large numbers of particles’ and not for ‘small numbers of particles’. A benefit of my proposal is that it avoids what can seemingly only be an arbitrary or gerrymandered resolution of this tension.

Secondly, we now understand that the atomic hypothesis ultimately requires a quantum interpretation. So, if there is to be an explanation of the success of classical statistical mechanics, it should come from a quantum mechanical description of the system and not from classical particle mechanics. That said, I reavow my lack of commitment to whether such a reductive explanation is possible, at least for the purposes of this paper. An interpretation of statistical mechanics, such as my proposal, does not require an explanation of the theory’s success. It is enough if it gives a picture of how the theory describes and explains the phenomena within its purview.

The picture my proposal provides may also strike some as puzzling for reasons having to do with that picture itself. For example, 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. This objection is easily deflected, since one need not make any realist commitment to phase space. In the first place, insofar as one thinks that realism is not an all or nothing affair with respect to the elements of a given formal framework, 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 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, so I supply this in the following section. Before turning to that, though, I think it is worthwhile to demonstrate that this burden of explaining away phase space is not really forced upon one advocating the ‘stochastic observables’ interpretation.

It is not a burden because it is possible to mold the above formalism somewhat so that the interpretation of observables as stochastic becomes even more manifest and phase space disappears. This formal modification is not at all necessary, 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 from above), 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 remove all reference to phase space and its microstates.

¹²I gratefully acknowledge two reviewers for pressing the various (serious) concerns to which I give my initial responses in the following paragraphs.

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 elements U in \mathcal{L} . Let ρ be the state of the system with respect to the measurable space (Γ, \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 . 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.

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:

$$\langle A \rangle_{\rho_A} = \int_{A[\Gamma]} \rho_A A \, dA. \quad (4)$$

The statistics predicted from formalizing statistical mechanics in this way, on probability spaces $\mathcal{A}[\Gamma]$ of observables, is 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.

The skeptical reader might object 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, such that the the claimed independence from phase space has not been shown. That is not really the case, though, since one could just as well work the other direction, namely, by choosing ρ_A and pulling this function back to some abstract space, for example phase space. Certainly the pullbacks are not going to be unique in general, but what is there to worry about? One just has to insure that the pullbacks of the set $\rho_{\mathcal{A}} = \{\rho_A, \rho_B, \dots\}$ agree appropriately in the target space (whatever that may be, phase space or otherwise). If so, then one has all the necessary resources to do statistical mechanics. Put this way, there is no obvious reason to demand an explanation for why statistical mechanical systems should be represented on something other than phase space. After all, should one demand such explanations 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 at all.”

4 Phase Space and Statistical Mechanics

Despite all I have said, I suspect many may 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.¹³ I certainly do grant that the phase space probability space is instrumentally useful, particularly 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, positing particular Hamiltonians, etc. I am happy to simply use phase space for these purposes while not taking it seriously representationally.

There are, however, two issues about phase space which still need addressing. 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, paradigmatically, a Hilbert space. Second, if phase space should not be taken too seriously for purposes of ontological interpretation, as I suggest, then why might phase

¹³“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 microstates” (Sklar, 1993, 291).

space be nonetheless a good way to represent statistical mechanical states? My answer 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. Once again, the phase space representation of quantum mechanics affords one precisely this possibility. Thus, there is at least one potential way to 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 point in phase space. Such points cannot combine to form objective 'superpositions' in the way needed to generate statistical mechanical states. Therefore, equilibrium statistical mechanics is naturally seen as a *more* general theory than classical mechanics, in that it can represent classical mechanical states probabilistically, as 'sharply-peaked distributions' (or 'pure states'), as well as a variety of other statistical states ('mixed 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.

Anyway, an essential point is that if one wishes to understand from where statistical mechanical states come (an external question, as it were), then one should look elsewhere for the provision of appropriate phase space representations. Classical mechanics cannot explain the requisite states. As a more fundamental, probabilistic theory, quantum mechanics is obviously an appealing place to look, and, as it happens, there is a relatively well-known, straightforward, and illuminating way to represent quantum states and operators on phase space (Kim and Noz, 1991; Case, 2008).

Let ρ be a density operator that represents a (mixed) quantum state, that is,

$$\rho = \sum_s p_s |\Psi_s\rangle\langle\Psi_s|, \quad (5)$$

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 $W_\rho : \Gamma \rightarrow \mathbf{R}$, which may be defined as follows:

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

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 states 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 in a similar way to quantum mechanical states (such transforms are called Weyl transforms), in particular as functions on phase space. Thus they may serve 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:

$$\langle A \rangle_\rho = \int_\Gamma W_\rho A d\Gamma. \quad (7)$$

Quantum mechanics is generally formulated in such a way that some Hilbert space is the state space of a system described by the theory. The state Ψ , for example, is an element of the 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. One is not led in this case to interpret quantum mechanical states as possessing classical microstates which correspond to phase space points merely in virtue of a phase space representation of quantum mechanics existing. Of course one may try to find classical microstates in this way. However, in typical observer-independent interpretations of quantum mechanics, the putative microstates would be treated instrumentally in quantum mechanics on phase space while the quantum state itself would be taken as real. I suggest that what is true here with quantum mechanics is true also with statistical mechanics. Just as phase space may be a useful representational structure in quantum mechanics for some incidental purposes, it is useful in statistical mechanics.

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. Rather, I mean to show that phase space can simply be a vehicle for representing underlying microphysical structure in an indirect way. Thus, one is not forced in any way to interpret the ‘fundamental’ ontology of the theory in line with the supposed classical particle states of phase space. There is no deep significance in 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.¹⁴ Accordingly, it is not necessary to reformulate statistical mechanics to avoid referring to phase space (as I did in §3). Phase space is surely 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.¹⁵

5 Alternative Foundations

In traditional interpretations of statistical mechanics some intra-theoretical foundational relationship is assumed between statistical mechanics and classical mechanics (Wallace, 2017). One common idea is to treat classical particle mechanics as if it were a fundamental theory—at least as 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.¹⁶ 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).

As I have been urging, the assumption of a necessary relationship between classical and statistical mechanics in such ways should strike us as somewhat odd: the world is, so far as we know, quantum mechanical and not classical mechanical.¹⁷ There is a multitude of possible inter-theoretic relations between

¹⁴“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).

¹⁵Here is a good point to note a relevant objection raised by a referee, the substance of which is that, as the Hamiltonian of a system governs the motion of the particles of the system, it gives a reason to endorse the reality of the particles it governs. On my proposal the Hamiltonian is not so understood; it instead governs the time evolution of the statistical state of a system and also serves as an observable giving the value of the energy of the system. Of course, Hamiltonians are usually defined on phase space. However, just as they induce a flow on phase space, they also induce a flow on the space of probability distributions associated to that phase space. So, by essentially the same argument that phase space need not be interpreted realistically, that is, as indicating particle content in the theory, I argue too that the Hamiltonian need not indicate it either. Consider too that the utility of Hamiltonians in quantum mechanics does not indicate that there are classical particles underlying the quantum wave function description.

¹⁶‘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).

¹⁷‘Why should we consider quantum issues when working in the foundations of statistical physics? The simple (too simple)

quantum mechanics, statistical mechanics, and classical mechanics to consider besides two simplistic reductions: statistical mechanics to classical mechanics, classical mechanics to quantum mechanics. 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 has also become unnecessary since the advent of quantum mechanics, as there are other potential ways of grounding the theory in one that is more fundamental.

A key virtue of 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 in classical particle mechanics. Besides the complaints mentioned in the introduction, Sklar finds this ‘revisionist’ approach to the ontology of statistical mechanics objectionable also because it fails to complete the historical project of statistical mechanics. For example, he claims that “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, denying the necessity of a particle ontology for statistical mechanics can help solve some crucial and puzzling questions, in particular by re-directing foundational efforts from a 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. Some popularly debated questions do fall by the wayside if one adopts my proposal, namely, those that depend on the traditional ontological suppositions. Yet 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 us with some questions. How can a *probability measure* represent the state of a system? How can this alternative interpretation coherently avoid a commitment to some particular foundational viewpoint? To some extent these questions have already been essentially answered by my discussion, but for lingering doubters I offer a couple of connections to recent philosophical work which may usefully illuminate the account further.

As urged especially by Callender, 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)

In this vein, 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 (including classical particles) in a variety of ways. In any case, it is essentially by treating statistical mechanics as a special science that the stochastic observables interpretation avoids an unnecessary commitment to a particular intrinsic foundation.

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).

Structuralism encourages the further view that it is unnecessary to posit any additional primitive ontology for the theory over and above the structure given by the theory itself. Structuralism, in its most benign form, should not be seen as an effort to ‘get rid of’ ontology altogether but rather as an effort to proof our metaphysical thinking of naive ontological prejudices.¹⁸ 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 both within physics (including, to some extent, quantum statistical mechanics) and within philosophy. However, we now know much more about microphysics than the founders of the theory, and a foundational account that presupposes their ontology should by now be seen 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 our philosophical and foundational endeavors.

Thus, 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 an unavoidable stochastic element which is naturally represented mathematically by probability measures. This is what it means, roughly, to take probability measures seriously as statistical mechanical states, with the probabilities representing the stochasticity of the system’s macroscopic observables. In much the same way that one can characterize quantum states as probability measures, one can also understand statistical states as probability measures (more easily, in fact, since classical observables commute). The metaphysical upshot of such probabilistic interpretations is that there is an interesting stochasticity in the world, at least at the level of description captured by those theories. The *revisionism* encapsulated by these interpretations is a profound metaphysical discovery, which should provoke further thought and investigation and not a retreat to a *reversionism* which attempts to preserve the traditional ontological by-gones.

I close by suggesting that the adoption of this interpretation opens up many new possibilities for foundational and philosophical work. I will mention only a couple; the reader perhaps can envision more. First, the possibilities of interpreting statistical mechanics sketched in §3 suggests interesting connections to the interpretation of quantum mechanics and the measurement problem. There are analogies between popular interpretations of quantum mechanics and these 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.¹⁹ Second, 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.²⁰

¹⁸“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).

¹⁹See (McCoy, 2018b) for an exploration of the interpretive analogies between statistical mechanics and quantum mechanics, including a discussion of how the measurement problem and primitive ontology figure into these analogies.

²⁰Such 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, D. (2000). *Time and Chance*. (Cambridge: Harvard University Press)
- Batterman, R. (1990). Irreversibility and Statistical Mechanics: A New Approach? *Philosophy of Science*, 57, 395–419
- Batterman, R. (1991). Randomness and Probability in Dynamical Theories: On the Proposals of the Prigogine School. *Philosophy of Science*, 58, 241–263
- Bishop, R. (2004). Nonequilibrium statistical mechanics Brussels-Austin style. *Studies in History and Philosophy of Modern Physics*, 35, 1–30
- Bradley, S. (2017). Are objective chances compatible with determinism? *Philosophy Compass*, 12, 1–11
- Bricmont, J. (1996). Science of Chaos or Chaos in Science? *Annals of the New York Academy of Science*, 775, 131–175
- Callender, C. (1997). What is “The Problem of the Direction of Time”? *Philosophy of Science*, 64 (Proceedings), S223–S234
- Callender, C. (1999). Reducing Thermodynamics to Statistical Mechanics: The Case of Entropy. *The Journal of Philosophy*, 96, 348–373
- Callender, C. (2001). Taking Thermodynamics Too Seriously. *Studies in History and Philosophy of Modern Physics*, 32, 539–553
- Callender, C. (2011). The Past Histories of Molecules. (In C. Beisbart, & S. Hartmann (Eds.), *Probabilities in Physics* (pp. 83–113). Oxford: Oxford University Press)
- Callender, C. & Cohen, J. (2010). Special Sciences, Conspiracy and the Better Best System Account of Lawhood. *Erkenntnis*, 73, 427–447
- Cao, T. Y. (2003). Structural Realism and the Interpretation of Quantum Field Theory. *Synthese*, 136, 3–24
- Case, W. (2008). Wigner functions and Weyl transforms for pedestrians. *American Journal of Physics*, 76, 937–946
- Clark, P. (2001). Statistical Mechanics and the Propensity Interpretation of Probability. (In J. Bricmont, D. Dürr, M. Galavotti, G. Ghirardi, F. Petruccione, & N. Zanghi (Eds.), *Chance in Physics* (pp. 271–281). Berlin: Springer Verlag)
- Ellis, R. (2006). *Entropy, Large Deviations, and Statistical Mechanics*. (Berlin: Springer-Verlag)
- Emch, G. (2005). Probabilistic issues in statistical mechanics. *Studies in History and Philosophy of Modern Physics*, 36, 303–322
- French, S. & Ladyman, J. (2003). Remodelling Structural Realism: Quantum Physics and the Metaphysics of Structure. *Synthese*, 136, 31–56
- Frigg, R. (2008). A Field Guide to Recent Work on the Foundations of Statistical Mechanics. (In D. Rickles (Ed.), *The Ashgate Companion to Contemporary Philosophy of Physics* (pp. 99–196). London: Ashgate)

- Guttman, Y. (1999). *The Concept of Probability in Statistical Physics*. (Cambridge: Cambridge University Press)
- Hemmo, M. & Shenker, O. (2012). *The Road to Maxwell's Demon*. (Cambridge: Cambridge University Press)
- Kadanoff, L. (2000). *Statistical Physics*. (Singapore: World Scientific)
- Kim, Y. & Noz, M. (1991). *Phase Space Picture of Quantum Mechanics*. (Singapore: World Scientific)
- Krylov, N. (1979). *Work on the Foundations of Statistical Physics*. (Princeton: Princeton University Press)
- Lavis, D. (2001). The Concept of Probability in Statistical Mechanics. (In B. Sidharth & M. Altaisky (Eds.), *Frontiers of Fundamental Physics 4* (pp. 293–308). Berlin: Springer Verlag)
- Lebowitz, J. (1999). Statistical Mechanics: A Selective Review of Two Central Issues. *Reviews of Modern Physics*, 71, S346–S357
- Loewer, B. (2001). Determinism and Chance. *Studies in History and Philosophy of Modern Physics*, 32, 609–620
- Ma, S-K. (1986). *Statistical Mechanics*. (Singapore: World Scientific)
- McCoy, C. (2018). No Chances in a Deterministic World <http://philsci-archive.pitt.edu/id/eprint/13166> (accessed 2018-05-01)
- McCoy, C. (2018). Interpretive Analogies Between Quantum and Statistical Mechanics. <http://philsci-archive.pitt.edu/id/eprint/14479> (accessed 2018-05-01)
- Meacham, C. (2010). Contemporary Approaches to Statistical Mechanical Probabilities: A Critical Commentary—Part I: The Indifference Approach. *Philosophy Compass*, 5, 1116–1126
- Myrvold, W. (2016). Probabilities in Statistical Mechanics. (In C. Hitchcock & A. Háyeek (Eds.), *Oxford Handbook of Probability and Philosophy* (pp. 573–600). Oxford: Oxford University Press)
- North, J. (2010). An empirical approach to symmetry and probability. *Studies in History and Philosophy of Modern Physics*, 41, 27–40
- Pathria, R. (1996). *Statistical Mechanics*. (Oxford: Butterworth-Heinemann)
- Penrose, O. (1970). *Foundations of Statistical Mechanics*. (Oxford: Pergamon Press)
- Prigogine, I. (1962). *Non-equilibrium Statistical Mechanics*. (New York: John Wiley & Sons)
- Prigogine, I. (1996). *The End of Certainty*. (New York: The Free Press)
- Sklar, L. (1979). “Probability as a Theoretical Concept.” *Synthese*, 40, 409–414
- Sklar, L. (1993). *Physics and Chance*. (Cambridge: Cambridge University Press)
- Uffink, J. (2007). Compendium of the Foundations of Classical Statistical Physics. (In J. Butterfield & J. Earman (Eds.), *Philosophy of Physics* (pp. 923–1074). Amsterdam: Elsevier)
- Uffink, J. (2011). Subjective Probability and Statistical Physics. (In C. Beisbart & S. Hartmann (Eds.), *Probabilities in Physics* (pp. 25–49). Oxford: Oxford University Press)

- van Lith, J. (2001). Ergodic Theory, Interpretations of Probability and the Foundations of Statistical Mechanics. (2001). *Studies in History and Philosophy of Modern Physics*, 32, 581–594
- von Plato, J. (1994). *Creating Modern Probability*. (Cambridge: Cambridge University Press)
- Wallace, D. (2001). Implications of quantum theory in the foundations of statistical mechanics. <http://philsci-archive.pitt.edu/id/eprint/410> (accessed 2018-05-01)
- Wallace, D. (2015). The quantitative content of statistical mechanics. *Studies in History and Philosophy of Modern Physics*, 52, 285–293
- Wallace, D. (2017). “Inferential versus Dynamical Conceptions of Physics.” (In O. Lombardi, S. Fortin, F. Holik, & C. López (Eds.), *What is Quantum Information?* (pp. 179–206). Cambridge: Cambridge University Press)
- Winsberg, E. (2008). Laws and chances in statistical mechanics. *Studies in History and Philosophy of Modern Physics*, 39, 872–888