Probabilities in Statistical Mechanics: Objective, Subjective, or a Bit of Both?

Wayne C. Myrvold
Department of Philosophy
The University of Western Ontario
wmyrvold@uwo.ca

May 31, 2011

Abstract

This paper addresses the question of how we should regard the probability distributions introduced into statistical mechanics. It will be argued that it is problematic to take them either as purely subjective credences, or as objective chances. I will propose a third alternative: they are *almost objective* probabilities, or *epistemic chances*. The definition of such probabilities involves an interweaving of epistemic and physical considerations, and so cannot be classified as either purely subjective or purely objective. This conception, it will be argued, resolves some of the puzzles associated with statistical mechanical probabilities; it explains how probabilistic posits introduced on the basis of incomplete knowledge can yield testable predictions, and it also bypasses the problem of disastrous retrodictions, that is, the fact the standard equilibrium measures yield high probability of the system being in equilibrium in the recent past, even when we know otherwise. As the problem does not arise on the conception of probabilities considered here, there is no need to invoke a Past Hypothesis as a special posit to avoid it.

Keywords: Statistical mechanics; thermodynamics; probability; chance; method of arbitrary functions.

the principles of statistical mechanics are to be regarded as allowing us to make reasonable predictions as to the future condition of a system, starting from incomplete knowledge of its initial state.

Richard C. Tolman (1938, p.1).

1 Introduction

Considerations of probability enter essentially in the formulation of statistical mechanics. There are, *prima facie*, two candidates for understanding such probabilities: subjective degrees of belief, and objective chance. Both seem problematic in the context of statistical mechanics. Probabilities are introduced because of incomplete knowledge of the state of a system—which suggests a subjective reading—but are nonetheless used to generate testable predictions, which suggests something objective. Furthermore, for an isolated system that starts out in disequilibrium and then relaxes to equilibrium, no probability-function over the full state of the system, whether it be regarded as a subjective or objective probability, evolves into a standard equilibrium distribution. One symptom of this is that the fine-grained Gibbs entropy of an isolated system does not increase. Another is the disastrous retrodictions that result from taking an equilibrium distribution seriously as a probability distribution over the microstates of the system: the standard equilibrium distributions yield high probability that the system was in equilibrium in the recent past, even when we know that it was not.

The key to resolving these puzzles lies in some work of Poincaré. Poincaré argued that, for certain systems, a wide range of probability distributions will be taken via the dynamics of the system into distributions that yield approximately the same probabilities for some statements about the macrostate of the system. If we restrict attention to a limited set of properties of the system, namely, those that can be revealed by feasible measurements, then there is good reason to believe that the equilibrium distributions yield approximately the same probabilities for the results of such measurements as would the result of evolving any probability distribution over initial conditions represented by a density function that doesn't vary too quickly over the phase space of the system. Such probability distributions are also those that can be taken to represent the beliefs of reasonable agents whose knowledge of the system is limited to the results of macroscopic measurements. We don't need to argue that *arbitrary* probability functions about initial conditions have this property, nor do we need to fix a unique probability function that all rational agents would be obliged to have, as the dynamics of the system can wash out considerable differences.

If it so happens that the dynamics of the situation wash out differences between credence-function of any two reasonable agents, with regards to the results of feasible experiments, we cannot come to this conclusion on the basis of either epistemic considerations or physical considerations alone. On the epistemic side, we require a restriction to a class of probability-functions that can represent the belief-states of reasonable agents with limited access to information about the system, and we require a limitation to certain sorts of measurements, since, for an isolated system, no dynamics will wash out *all* differences between probability functions. On the physical side, we need the right sort of dynamics. If it is the case that the dynamics lead to probability distributions that yield approximately the same probabilities for the results of certain measurements, given as input any reasonable credence-function, then this fact is a matter of physics, as is the values of such probabilities. We will call the probabilities that emerge in such a scenario *almost objective* probabilities, or *epistemic chances*.[1]

All of this will be spelled out in more detail, below. It will be argued that the probabilities that emerge from the mix of epistemic and physical consideration just outlined are the right sort to play the role required of them in statistical mechanics.

For ease of exposition, we will be primarily concerned in this paper with classical statistical mechanics. Many of the conceptual issues will be essentially the same for quantum statistical mechanics; in particular, if a quantum mechanical mixed state is introduced because of imperfect knowledge of the actual state of the system, then the status of such mixtures will be much the same as that of classical mixed states. It should be noted, however, that, in light of the puzzles associated with classical statistical mechanical probabilities, some authors have suggested that quantum mechanics is required to make sense of statistical mechanics. In this vein, David Albert (2000, pp. 154–162) takes recourse to state-vector collapse, regarded as a real, and chancy, physical process. Along different lines, Linden et al. (2009) suggest that it is necessary to consider the reduced state of a system that is entangled with its environment. Both of these proposals merit serious consideration as an account of the origin of statistical mechanical probabilities, consideration which is beyond the scope of the current paper. It is, however, an aim of this paper to argue that quantum mechanics is not *needed* in order to make sense of statistical mechanical probabilities; the puzzles can be resolved within the classical context, and classical statistical mechanics can stand on its own.

2 Concepts of Probability

2.1 Two senses of "probability"

The word "probability" is used in at least two distinct senses. One sense, which historically is the older,[2] has to do with degrees of belief. There is another sense on which probability is associated, not with agents and their beliefs, but with physical set-ups such as the roll of a pair of dice. The probability of getting a pair of sixes on a roll of the dice is spoken of as being a characteristic of the dice and the circumstances of the throw—the "chance

[1] This term is borrowed from Schaffer (2007).

[2] See Hacking (1975) for a masterful overview of the history.

set-up"—whether or not this probability is known to anyone (for example, unbeknownst to all, the dice might be biased). The two have much in common. For one thing, there is a formal correspondence, as an assignment of objective chances and the degrees of belief of an ideally rational agent will both satisfy the usual axioms of probability. But they are distinct, nonetheless.

The literature on interpretations of probability sometimes suggests that we have to choose one or the other of these senses, or some other sense, as *the* meaning of probability. The attitude adopted in this paper is that both of these senses are useful concepts, and we would be foolish to discard either. What is important is that we not conflate the two. In what follows, we will use the word "chance" for objective probabilities, "credence" when we are speaking of the a degree of belief of some (perhaps idealized) agent.

Inattention to the distinction between objective chance and subjective credence carries a risk of falling into an illusion that ignorance can, by some sort of alchemy, be transformed into knowledge. Suppose that we know of no reason why a coin will land heads rather than tails, and accordingly assign these alternatives equal probability. Since there is an equal probability of heads and tails, we then predict, with a high degree of confidence, that a large number of tosses will yield roughly equal heads and tails. Ignorance has yielded knowledge.

Distinguishing between chance and credence, it is easy to see what has gone wrong. To have one's credence equally divided between heads and tails is not the same as being certain that the chances of heads and tails are equal. Equal credence for heads and tails is compatible with a wide variety of credences about the chance of heads, as any distribution of credences about the chance of heads that is symmetrical around $1/2$ yields equal credence for heads and tails. One might, for example, have credences about the chances represented by a flat density function. That this is different from being certain that the chance is $1/2$ can be seen from the fact that different credences are yielded for the results of multiple tosses. An agent whose credences about the chance of heads are flat has credences $\{1/3, 1/6, 1/6, 1/3\}$ in the possible outcomes $\{HH, HT, TH, TT\}$ of a pair of tosses, whereas an agent who is certain that the coin-toss is fair has equal credence $1/4$ in each of these possible outcomes.[3]

2.2 Learning about chances

The chance of heads on a coin toss, if it is regarded as an objective feature of the set-up, is *ipso facto* the sort of thing that we can have beliefs about, beliefs that may be correct or incorrect, better or worse informed. Under certain conditions, we can learn about the values of chances.

Particularly conducive to learning about chances are cases in which we have available (or can create) a series of events that we take to be similar in all aspects relevant to their chances, that are, moreover, independent of each other, in the sense that occurrence of one

[3]In general, as can be shown by a simple calculation, a flat distribution over chances of heads makes every possibility for the number of heads in N tosses equally probable.

does not affect the chance of the others. The paradigm cases are the occurrence of heads on multiple tosses of the same coin, occurrence of a six on multiple throws of the same die, and the like. Consider a sequence of N coin tosses. If, on each toss, the chance of heads is λ, then the chance of obtaining any particular sequence of results having m heads and $N-m$ tails is

$$\lambda^m (1-\lambda)^{N-m}.$$

Considered as a function of λ, this is peaked at the observed relative frequency m/N, and becomes more sharply peaked, as N is increased.

Let E be the proposition that expresses the sequence of results of these N tosses. Consider an agent who has some prior credences about the chance of heads, and updates them by Bayesian conditionalization. Let these credences be represented by a density function $f(\lambda)$; that is, our agent's credence that the chance is in an interval Δ is given by

$$cr(\lambda \in \Delta) = \int_\Delta f(\lambda)\, d\lambda.$$

We also define a likelihood function $l(E|\lambda)$ that satisfies

$$cr(E \,\&\, \lambda \in \Delta) = \int_\Delta l(E|\lambda) f(\lambda)\, d\lambda.$$

Updating by Bayesian conditionalization on the evidence shifts the credence-density function:

$$f(\lambda) \;\Rightarrow\; f(\lambda|E) = \frac{l(E|\lambda)\, f(\lambda)}{cr(E)}.$$

It seems natural to suppose—and, indeed, in the statistical literature this is typically assumed without explicit mention—that our agent's credences set $l(E|\lambda)$ equal to the chance of E according to the hypothesis that the chance of heads on each toss is λ, as is required by what Lewis (1980) has dubbed the *Principal Principle*.[4] This gives

$$f(\lambda) \;\Rightarrow\; f(\lambda|E) = \frac{\lambda^m (1-\lambda)^{N-m}}{cr(E)}\, f(\lambda).$$

This has the consequence that, provided our agent's prior credences don't assign credence zero to some interval containing the observed relative frequency, her credence in chance-values close to the observed relative frequency is boosted, and her credence in other values, diminished. Moreover, since the likelihood function $\lambda^m(1-\lambda)^{N-m}$ is more sharply peaked, the larger the number of trials, relative frequency data becomes more valuable for narrowing credence about chances as the number of trials is increased.

[4]Since one of the chief sources of significance of the Principal Principle is the role it plays in learning about chances, readers are urged to resist the temptation to gloss the Principle as the injunction to set one's credence equal to the chance, when the chance is known. It is when the chance is *not* known that the Principal Principle is most valuable!

Note that there are three distinct concepts at play here: chance, credence, and relative frequency in repeated trials. None of these three is to be identified with any of the others. They do, however, connect in a significant way: relative frequency data furnish evidence on which we update credences about chances.

2.3 Two non-senses of "probability"

Missing from our classification are the two conceptions of probability that arise most frequently in discussions of probability by physicists. These are the classical conception of probability, founded on a Principle of Indifference, and frequentism. A few words are in order about why neither of these are adequate.[5]

In the classical vein, there is a temptation to attempt to define the probability of an event A as the ratio of the number of possible cases favourable to A to the total number of possible cases, if the state space is discrete, or as the ratio of the volume of the state space in which A holds to the total available volume of state space, in the continuum case. This can't be the whole story, however. In the discrete case, a judgment is needed as to which way of partitioning the state space yields equiprobable cases. In the continuum case, a judgment is required about which measure on the space is the appropriate one. We may say: a uniform measure, represented by a flat density function, but then we must remind ourselves that a probability distribution that is uniform with respect to one parameterization of the space will not be uniform with respect to others. If the appropriate distribution is uniform, then a judgment is required about which variables it is uniform in.

One of the best comments on attempts to define probability in terms of mere counting of cases occurs in Laplace's *Philosophical Essay on Probabilities*, in which an incautious formulation is first enunciated, then corrected.

> *First Principle.*— The first of these principles is the definition itself of probability, which, as has been seen, is the ratio of the number of favorable cases to that of all the cases possible.
>
> *Second Principle.*— But that supposes the various cases equally possible. If they are not so, we will determine first their respective possibilities, whose exact appreciation is one of the most delicate points of the theory of chance (Laplace, 1951, p. 11).

It is certainly true that, given a judgment that a certain partition of the state space is to be regarded as equiprobable, such a judgment fixes the probability of all boolean combinations of elements of this partition; moreover, this is a very useful fact, as it reduces a great many problems in the theory of probability to combinatorics. Laplace's First Principle does not

[5]The literature on both of these topics is vast, and some readers may find the following inadequate. But a thorough discussion of these matters is beyond the scope of this paper.

suffice as a *definition* of probability, since, as pointed out by Laplace himself, it requires supplementation by a judgment of which cases are equiprobable.[6]

There are also attempts to identify objective probabilities with relative frequencies of events, in either actual or hypothetical ensembles of events. There is, indeed, a connection between chances and relative frequencies. For example, if we consider the case of a ball being drawn from an urn containing a number of balls, with each ball having an equal chance of being drawn, then the chance that the drawn ball will be black is equal to the proportion, or relative frequency, of black balls in the urn. We cannot, however, simply equate the chance of a black ball being drawn with this proportion, as the conclusion that the chance of drawing a black ball is equal to the proportion of black balls relies on the condition that each ball have an equal chance of being drawn, and this requires a notion of chance distinct from that of relative frequency of balls in the urn. The situation is worse for other events. In the case of drawing from an urn, it is clear that the relevant reference class should be the balls in the urn. In other cases it is less clear what the appropriate reference class should be, and what chance we ascribe to an event may vary widely depending on the reference class.

Recourse might be made to limiting relative frequencies in a hypothetical infinite sequence of repeated events of the same type; we may be tempted to define the chance of heads in a coin toss as the value that the relative frequency would converge to, if the coin were to be tossed infinitely many times. But why should we think that there is such a value? How are we to evaluate the counterfactual? Appeal may be made to the Strong Law of Large Numbers, which assures us that, in an infinite sequence of identically distributed and independent events, the relative frequency of any outcome-type will converge to a limiting value. But we need to be careful. There are, of course, possible sequences on which the relative frequency does not converge, and possible sequences on which the relative frequency converges to the wrong value. What the Strong Law says is that the set of such sequences has probability zero. This requires us to to be able to ascribe probabilities to propositions regarding whether or not there will be a limiting relative frequency, and to propositions regarding the value of the limiting relative frequency, if there is one, and this requires a notion of probability distinct from limiting relative frequency. Although relative frequencies of events in repeated trials have a bearing on chances, in that they are in many cases our best evidence about the values of these chances, they are conceptually distinct from chances.

Both the classical conception and the frequency conception are attempts to introduce an objective notion of probability that is compatible with deterministic laws of nature. We will be able to sidestep the issue, as it will be argued below that, whether or not there is a fully objective notion of probability available, we can introduce a notion that will suffice to play the role, for the purposes of statistical mechanics, that objective probability is meant to play.

[6] "Equally possible" is Laplace's phrase. For discussion of the meaning of this, see Hacking (1971). Laplace has been accused of circularity, on the grounds that "equally possible" can only mean *equally probable*. Even if Laplace can be defended against the charge, the point remains: what we have does not suffice as a definition of probability unless supplemented by an account of which cases are to be regarded as equipossible.

3 Statistical Mechanical Probabilities: Epistemic, or Objective?

3.1 Puzzles about statistical mechanical probabilities

The state of a classical system is represented by a point in its phase space, which is specified by specifying the values of all coordinates q_i in its configuration space and their conjugate momenta p_i. These change over time according to *Hamilton's equations*:

$$\dot{q}_i = \frac{\partial H}{\partial p_i}, \quad \dot{p}_i = -\frac{\partial H}{\partial q_i},$$

where H is the Hamiltonian of the system, that is, the total energy of the system, expressed as a function of coordinates and momenta $\{q_i, p_i\}$. We will use x as a variable ranging over phase space points: that is, over full specifications of all coordinates and momenta. As the state of the system evolves, so too will any probability distribution over the phase space. If $\rho(x,t)$ is a probability density function over the state of the system at time t, then *Liouville's equation* expresses the time-dependence of this probability density function:

$$\frac{\partial \rho}{\partial t} + \sum_i \left(\frac{\partial \rho}{\partial q_i}\frac{\partial H}{\partial p_i} - \frac{\partial \rho}{\partial p_i}\frac{\partial H}{\partial q_i} \right) = 0.$$

Any distribution that is uniform with respect to canonical phase space variables will be unchanging in time, as will be any distribution given by a density function that is a function only of the total energy H.

Statistical mechanics invokes certain standard probability distributions for systems in equilibrium. For an isolated system with conserved total energy known to lie in a small interval $E, E + \Delta E$, we use the *microcanonical distribution*, on which regions of phase space with equal phase-space volume have equal probability, provided that the energies associated with points in the region lie entirely in the interval (one can also consider a projection of this distribution down to the surface of energy E). For a system that is in contact with a heat bath of temperature T, one uses the canonical distribution, with density function

$$\rho(x) = Z^{-1} e^{-H(x)/kT},$$

where k is Boltzmann's constant, and Z is chosen so that the function is normalized. Given a phase-space probability density ρ, we define the *Gibbs entropy*

$$S[\rho] = -k \int \rho(x) \log \rho(x)\, dx,$$

where the integral is taken over the entire phase space of the system. It follows from Liouville's equation that the Gibbs entropy of an isolated system is constant.

The question we want to ask is: are these probabilities to be regarded as subjective degrees of belief, or objective chances?

Textbook expositions typically begin by observing that the systems to which we apply thermodynamics possess a large number of degrees of freedom, and that our knowledge of the system usually consists of measured values of a relatively small number of parameters. The state of the system is, therefore, incompletely known. For this reason, we introduce a probability distribution over the states of the system compatible with what is known about it. One textbook tradition, with its roots in Tolman (1938), bases probability assignments in statistical mechanics on what Tolman calls the "hypothesis of equal *a priori* probabilities in the phase space." About this postulate, Tolman writes,

> Although we shall endeavour to show the reasonable character of this hypothesis, it must nevertheless be regarded as a postulate which can ultimately be justified only by the correspondence between the conclusions which it permits and the regularities in the behaviour of actual systems which are empirically found (Tolman, 1938, p. 59).

Tolman's argument for the reasonableness of adopting a probability distribution that is uniform with respect to canonical variables is based on Liouville's theorem, which demonstrates that such a distribution is a stationary distribution; this shows that "the principles of mechanics do note themselves include any tendency for phase points to concentrate in particular regions of the phase space" (p. 61).

> Under the circumstances we then have no justification for proceeding in any manner other than that of assigning equal probabilities for a system to be in different equal regions of the phase space that correspond, to the same degree, with what knowledge we do have as to the actual state of the system. And, as already intimated, we shall, of course, find that the results which can then be calculated as to the properties and behaviour of systems do agree with empirical findings (p. 61).

This suggests an application of a Principle of Indifference, albeit not an incautious one that disregards the need for a judgment about the variables with respect to which the distribution is to be uniform. We seem to be skirting dangerously close to the alchemy of turning base ignorance into golden knowledge. It looks as if ignorance probabilities, introduced on the basis of a Principle of Indifference, yield empirical predictions which are found to be corroborated. Indeed, this is strongly suggested by some subsequent textbooks. To take one example, in the opening chapter of E. Atlee Jackson's *Equilibrium Statistical Mechanics*, a version of the Principle of Indifference is introduced.

> to *predict* the probability of a certain event, one uses the general rule:

> **If there is no apparent reason for one event to occur more frequently than another, then their respective probabilities are assumed to be equal** (Jackson, 1968, p. 8).

On the basis of this and some plausibility considerations, Jackson introduces

> **The Basic Assumption of Statistical Mechanics**. All microstates of a system that have the same energy are assumed to be equally probable.

This is immediately followed by the remark,

> This simple *assumption* is the basis of all of statistical mechanics. Whether or not it is valid is a matter that can be only be settled by comparing the predictions of statistical mechanics with actual experiments. To date there is no evidence that this basic assumption is incorrect. A little thought shows that this agreement is indeed remarkable, for our basic assumption is little more than a reflection of our total ignorance about what is going on in the system (p. 83).

It would, indeed, be remarkable if an assumption of ignorance could be turned into reliable empirical predictions. But ignorance cannot be transformed into knowledge. There must be something else going on.

We may be tempted to simply reject all of this as confusion stemming from a conflation of the objective and subjective senses of probability. If we do so, then, it seems as if we must adopt a single reading of "probability": either objective or subjective. Neither one seems able to do all that we ask of it, however. Probabilities are introduced in the first place because we have incomplete knowledge of a system. Yet on the basis of these probabilities we calculate expectation values of measurable quantities, as well as probability distributions for deviations from these expectation values, and we find our expectations to be confirmed by experiment. This suggests that the probabilities in question are something more than a reflection of ignorance. We have a puzzle: how can it be that a postulate about probabilities, introduced because of incomplete knowledge of the state of the system, can be tested by experiment?

Furthermore, it seems problematic to take the standard statistical mechanical probability distributions as representing *either* our credences about the system's microstate, or objective chances regarding the system's microstate, as, for both interpretations, we run up against the familiar problem of disastrous retrodictions. To see this, consider the example of free expansion of a gas. Suppose that at time t_0, a gas is confined, and known to be confined, by a partition to one compartment of a box. It matters not for the purposes of this example whether, prior to t_0, the system is isolated and its energy known, so that standard statistical mechanics recommends a microcanonical distribution, or it is in contact with a heat bath of known temperature, so that we use the canonical distribution. Let the partition be removed, and the gas permitted to expand into the remainder of the accessible volume, while remaining adiabatically isolated. Let t_1 be some time sufficiently long after the removal of the partition for the system to have relaxed to its new equilibrium.

Let ρ_0 be the phase-space distribution we use for the system's state at time t_0. Let ρ_1 be the result of evolving, via the dynamics of the system, the distribution ρ_0 to the time t_1. Let ρ_1^S be the standard equilibrium distribution assigned to the gas at t_1, that is, the microcanonical or canonical distribution over the over the region of phase space accessible to the gas at time t_1.

It should be noted, first, that ρ_1 is not the same as ρ_1^S. In some respects, they are very different. Let A_0 be the set of states in which the gas is entirely in the original compartment, and let A_1 the result of evolving A_0 from t_0 to t_1 Since the evolution of an isolated system preserves phase-space volume, A_1 has the same phase-space volume as A_0. A_1 is a minuscule fraction of the set of states available to the system at t_1. Therefore, ρ_1^S assigns miniscule probability to A_1, though ρ_1 assigns probability one. Furthermore, ρ_1 and ρ_0 have the same value of the Gibbs entropy, whereas the entropy of ρ_1^S is greater than that of ρ_0.

Since, at t_1, the system has relaxed to equilibrium, ρ_1^S is the probability distribution that will be used in statistical mechanics. What is its status? Does it represent our beliefs about the system, objective chances, or something else?

To take ρ_1^S as representing our beliefs about the state of the system at time runs up (at least, *prima facie*; see §3.5 for a qualification) against the problem of disastrous retrodictions. ρ_1^S assigns vanishingly small probability to the gas having been confined to its original compartment at time t_0, yet, by hypothesis, we know that this was the case. Nor does the probability assigned to a region of phase space by ρ_1^S represent an objective chance that the system's state lies in that region. As the system was confined to its original compartment at t_0, it has no chance of ending up at time t_1 at any point in phase space that doesn't back-evolve into that a state like that, though ρ_1^S assigns probability close to unity to the set of such points. Thus, it seems that the equilibrium probability distribution, applied at t_1, represents neither our knowledge of the system (as it discards the information that the system was not always spread out over the entire volume), nor objective chances about the state of the system.

3.2 Puzzles resolved: epistemic chances, or almost-objective probabilities

If the standard probability distributions represent neither objective chances nor our credences, why, then, do we bother with them?

Imagine a Laplacean demon, with perfect knowledge of the dynamics of some isolated thermodynamic system, and unlimited capacity to perform calculations on the basis of these dynamics. Consider once again our gas, initially confined to a subcompartment and then allowed to expand freely. Consider Bob, an agent with finite cognitive capacities, who is able to ask questions of the demon. If Bob tells the demon the credence-function representing his beliefs about the state of the system at time t_0, the demon will be able to evolve this probability distribution to time t_1. Knowing the result of such a calculation would tell Bob what his credences about the state of the gas at t_1 ought to be, given his credences about its state at t_0 and taking full advantage of the demon's expertise in classical mechanics. It is not at all clear, however, that the demon would be able to communicate this distribution, in all its detail, to Bob, or that it would be of much use to Bob if he had it, as the distribution would be monstrously complex; for one thing, the support of the distribution would be a complicated, finely fibrillated shape threading through the system's phase space. If, however, Bob's interests are in forming expectations about the results of feasible measurements (we might imagine him making bets on the outcomes of such measurements), then the demon's communication task is much simpler. For most practical purposes, the demon could simply provide Bob with the expectation value and dispersion of a finite number of measurable quantities.

It is conceivable—and, in fact, we have good reason to believe that this is the case for the sorts of systems to which we successfully apply thermodynamics—that the demon's communication task will be even simpler than this. Even though the time-evolute of Bob's credences about the gas at t_0 is enormously complex, if the dynamics are of the right sort there will be simplicity in this complexity, because the support in phase space of Bob's credence about initial conditions will end up so finely spread through phase space that the probabilities of results of macroscopically feasible measurements will not differ appreciably from those yielded by an equilibrium distribution that the demon can communicate succinctly, and Bob will be able to calculate with that.

If this is right, then it will also be true for agents whose credences about the state of the system at t_0 differ from Bob's, including some that differ widely. As the equilibrium distribution bears no trace of Bob's belief that the system was confined to one compartment at t_0, the demon will be able to give the same advice to another agent, Alice, who believes that at t_0 the gas was confined to the *other* compartment. Although the time-evolutes of Alice's and Bob's credences will be, in one sense, as different as they were before—they will still have disjoint supports—these differences will have washed out, as far as probabilities about the outcomes of macroscopic measurements are concerned.

When we have a situation like this, in which the dynamics of a system turn widely differing probability distributions over states of affairs at some time t_0 into distributions that yield virtually the same probabilities regarding certain coarse-grained propositions about states of affairs at a later time t_1, these probabilities are what have been called *almost objective* (Machina, 2004). Poincaré (1912, §92) pioneered the study of such probabilities, in connection with simple roulette-like systems; see von Plato (1983) for a masterful overview of the history, and (author's paper, forthcoming) for discussion. Poincaré called the method he used to analyze such set-ups the *method of arbitrary functions*.

To get probabilities about final conditions that are the same for arbitrary smooth probability densities over initial conditions, Poincaré had to pass to an unphysical limit. Specifically, he considered the limit in which the number of alternating red and black sectors of his wheel goes to infinity. For actual physical systems, it will not be the case that arbitrary credences about initial conditions, or even arbitrary credences with the same support and which are yielded by smooth probability densities, will yield approximately the same probability for measurements performed at t_1. One could, for example, cook up a probability function over the state of the system at t_0 that, instead of diffusing throughout the container, kept the gas within the original compartment.

Suppose, now, that Alice claimed that her credences about the gas at t_0 were concentrated on the class of states that yield anti-thermodynamic behaviour. It can be argued that, for any state in this class, miniscule perturbations of the state will suffice to take the state into one that yields the thermodynamic behaviour that we ordinarily expect. Alice's credence-density function would have to vary strongly over small distances in phase space. She would, therefore, be claiming very detailed knowledge about the state of the system. If her sources of information about the state of the system is limited to results of macroscopic measurements, we would rightly judge her credences to be unreasonable, even if they are coherent; she does not have knowledge of the system that is that detailed.

Though we will not attempt to define unique credences that any rational agent would be obliged to have, given a body of information, we will take it that there are some credences that are reasonable, given what an agent knows, and some that are not (there will, of course, be some vagueness about this). We will not need to suppose that the dynamics washes out difference between completely arbitrary credence-functions; for our purposes, it suffices that it does so for all reasonable credence functions.

Generalizing, suppose we have:

- a class $\mathcal{C}$ of credence-functions about states of affairs at time t_0 that is the class of credences that a reasonable agent could have, in light of information that is accessible to the agent,
- a dynamical map $\mathcal{T}$ that maps microstates at time t_0 to microstates at time t_1, which induces a map of probability distributions over states at time t_0 to distributions over states at time t_1.
- a set $\mathcal{A}$ of propositions about states of affairs at time t_1, whose truth-values can be ascertained by observation or experiment.

If, for some $A \in \mathcal{A}$, $\mathcal{T}\rho(A)$ has approximately the same value for all $\rho \in \mathcal{C}$, then we will call this common value an *almost-objective probability*, or the *epistemic chance* of A. Note that these probabilities are those yielded by evolving credences about the initial state via the actual dynamics of the system, whether these are known to the agent or not. We are not talking about intersubjective agreement; the value of an epistemic chance might be unknown to all, and might not represent the credence of any agent. Under propitious circumstances, however, agents can gather evidence about the values of such epistemic chances and by conditionalizing on such evidence come to agreement (see §3.4, below).

Quantities such as these are suited to play a role analogous to objective chance, even if the underlying dynamics are deterministic. Note that the definition includes both physical considerations, in that the dynamics must be the right sort, and epistemic considerations, having to do with limitations on accessible knowledge about the system. Whether or not a given proposition A will have an epistemic chance depends both on the dynamical map $\mathcal{T}$ and on the class of $\mathcal{C}$ of reasonable credences (though, if all goes well, will not depend too sensitively on the latter); the value of the epistemic chance of A, if there is such a value, is largely a matter of the dynamics.

It is sometimes said that the fundamental assumption underlying statistical mechanics is that the dynamics be ergodic, or mixing, or satisfy some other condition.[7] We do, indeed, require some condition on the dynamics, but ergodicity is not a necessary requirement. The conditions under which we speak of almost-objective probabilities may be met in connection with a system whose phase space consists of two invariant subsets of nonzero measure, provided that these are so finely intertwined that they cannot be distinguished by any feasible measurement. What is needed is something like the following.

> *Hypothesis of Washing-Out of Credence about the Past.* If a system is out of thermodynamic equilibrium at time t_0, and has relaxed to equilibrium by time t_1, then any distribution that results from applying the dynamical evolution of the system to some reasonable credence-function about the state of the system at time t_0 yields approximately the same probability for the results of feasible measurements performed after t_1.

It is sometimes said that the fundamental assumption underlying statistical mechanics must take the form of some assumption about probabilities, based on a version of the principle of indifference.[8] The difference between our approach and one based on such a principle is that no attempt is made to uniquely specify a probability distribution over microstates; we can accept different agents with the same knowledge of macroscopically accessible parameters disagreeing on probabilities concerning microstates at some time; provided that neither credence function is wildly unreasonable, these differences will soon wash out.

A common approach to statistical mechanics invokes coarse-graining. Instead of a probability distribution over the precise microstate of the system, one adopts some coarse-graining procedure to yield a probability distribution over coarse-grained states. A probability distribution over coarse-grained states yields probabilities only for observables that are expressible in terms of the coarse-grained description. From the perspective adopted here, this can be seen as one way of specifying a limited class of observables $\mathcal{A}$, about which there will be almost-objective probabilities.

The equilibrium distribution, applied at t_1. is being used as a *surrogate* for the time-evolute of the agent's credence about states of affairs at t_0, and it is so used for the limited purpose of providing probabilities for results of experiments performed after t_1. There is no

[7]See Berkovitz et al. (2006) for an insightful discussion of such conditions.

[8]Of particular relevance here is the approach of E.T. Jaynes, who bases statistical mechanics on a Principle of Maximum Entropy.

rationale for using it for earlier events; in particular, though it yields high probability that the system was in equilibrium at t_0, there is nothing paradoxical about this. We ought not to use this distribution for earlier events, as the very reason it was introduced is that our agent's knowledge of the past of the system has become largely irrelevant to results of future measurements. We thus don't encounter the sorts of disastrous retrodictions that lead David Albert (2000) to introduce his "Past Hypothesis."

3.3 The role of stationary distributions

Consider, again, our stock example of the free expansion of a gas. Suppose that, if we wait long enough after the partition is removed, differences between reasonable credences about states of affairs before the removal of the partition wash out, and there is an almost-objective chance for the result of any feasible measurement performed on the system at time t_1. Suppose, further, that no measurement performed after t_1 will permit us to determine how long the system has been in equilibrium; that is, the time-evolved credences of Charles, who takes the system to have been at equilibrium at t_0, will yield effectively the same probabilities for outcomes of measurements performed after t_1 as the time-evolved credences of Bob and Alice, who take the system to have been out of equilibrium at t_0. Then a stationary distribution will suffice as a surrogate for the evolute of Charles's credences about the state at t_0, and, since, by supposition, Charles' time-evolved credences yield effectively the same probabilities for results of measurements performed after t_1, as a surrogate for Alice's and Bob's time-evolved credences, as well. That is, if the evolution of the system washes out differences between Alice, Bob, and Charles, so that there is a distribution that will serve as common surrogate for the time-evolved credences of all three, then among distributions that suffice for this purpose will be a stationary distribution. Hence, in the quest for probability distributions that can yield almost-objective probabilities, stationary distributions will play a prominent role.

3.4 Testing hypotheses about almost-objective probabilities

For real systems having a macroscopic number of degrees of freedom, we may not know the exact Hamiltonian, and, even if we did, typically could not do the detailed calculation required to evolve a non-equilibrium credence-function over a substantial period of time. An agent might, nevertheless, believe that, for some propositions about the system, there exist almost-objective probabilities, even if she doesn't know the values of these probabilities. In such a case, our agent can entertain hypotheses about what these values are, and have degrees of belief in such hypotheses. We can imagine her wondering what advice would be given by a Laplacean oracle that knew the dynamics of the system in question and could evolve her credences about states of of affairs at t_0 into credences about states of affairs at t_1, and she can have credences about what that advice would be.[9] That is, epistemic chances, like objective chances, are the sorts of things that our agent can have credences about. Let A be some proposition about the system, and let our agent's credences about the value λ of the epistemic chance of A be given by a credence-density function $f(\lambda)$. She will also have credences of the form $cr(A \,\&\, \lambda \in \Delta)$, and these can be used to define a likelihood function $l(A|\lambda)$. Suppose, now, that her credence in these things is such that her credence in A, conditional on the supposition that the oracle would recommend credence λ in A, is equal to λ. That is,

$$l(A|\lambda) = \lambda.$$

This condition is an analog of the Principal Principle, applied to epistemic chances; call it the PPE. It is reasonable to expect our agent's credences to satisfy this. If her credences about state of affairs at time t_1 are not the evolutes of her credences about states of affairs at t_0, this is either a failure of coherence or reflects uncertainty (due to uncertainty about the actual dynamics, or what the result of applying those dynamics would be) about what would result from evolving her credences about the state at t_0. We can gloss the PPE as saying that, if our agent had access to a Laplacean oracle, learning consisting of conditionalization on its pronouncements would result in acceptance of its recommendations.

If A is a proposition whose truth or falsity can be ascertained by measurement or observation, and we have multiple copies of the system, then, provide her credences satisfy the PPE, Alice can do experiments and update her credences about what the recommendations of the oracle would be, in a manner precisely analogous to learning about chances, as outlined in §2.2. This will have the consequence that credence in hypotheses that accord higher epistemic chance to the observed results will be boosted relative to credences that accord them lower epistemic chance. With sufficient evidence, we can end up with arbitrarily high confidence that we know what the actual epistemic chances are.

Thus, though almost-objective probabilities have an epistemic aspect to them, they nonetheless are testable by experiment. We can, on the basis of the sorts of plausibility grounds mentioned by Tolman, conjecture that the standard distributions yield the correct epistemic chances, and then test this conjecture by experiment. The seeming paradox, that a postulate about probabilities introduced on the basis of ignorance can have testable consequences, is resolved.

3.5 Non-Liouvillean evolution and entropy increase

So far we have been talking about isolated systems. Isolation is at best approximate. If we begin with a probability distribution over system + environment and evolve it, then, though the evolution of the probability distribution over the state of system + environment

[9] This is a calculating oracle. We don't imagine that the oracle knows what the actual state of the system is; rather, it tells one what credences about states of affairs at time t_1 are consistent with one's credences about states of affairs at time t_0 and the laws of physics.

is measure-preserving, the marginal distribution of the system (that is, the distribution that results from restricting one's attention to the degrees of freedom of the system and averaging out the degrees of freedom of the environment) may increase in entropy.

The system will typically be small (few degrees of freedom) relative to the environment with which it interacts. There are a number of arguments, of varying rigor and generality, that show that, if the system interacts only weakly with its environment, so that the total Hamiltonian can be approximated by the sum of the system's energy and the environment's, then, provided the system's degrees of freedom are few compared to the environment's, then the marginal distribution for the system that is yielded by a microcanonical distribution over system + environment closely approximates a canonical distribution.[10] The result holds for both classical and quantum mechanics, and is often presented as a justification for using the canonical distribution for a system in thermal contact with a heat bath. Recently, this result has been generalized in interesting and significant ways. Although the usual quantum argument assumes a uniform distribution (a totally mixed state) for system + environment on its energy subspace, it has been shown that that this assumption can be considerably weakened. When the environment is much larger (that is, has a much higher dimensional Hilbert space) than the system, then most pure states of the system + environment yield reduced states for the system that closely approximate a canonical distribution.[11] This has the consequence that any probability distribution over the state of the system + environment will yield, as marginal for the system, an approximation to a canonical distribution, except for those probability distributions that give large weight to the exceptional pure states that yield marginals for the system that aren't close to canonical.

This result has a classical analog. It can be shown that most (in some sense of 'most') classical probability distributions over a composite of a system weakly coupled to a much larger environment yield marginals for the system that approximate the canonical distribution (Plastino and Daffertshofer, 2008). This result was proven for a discrete classical state space; it would be interesting to see an extension of the result to distributions over a continuous classical phase space.

The upshot of all this is that, even if our agent's credences about the state of a system and its environment differ wildly from the standard statistical mechanical probability distributions, her credences about the system itself might be closely approximated by the canonical distribution, which can then be used a as a proxy for her actual credences when calculating probabilities of the results of subsequent measurements.

Thus, for non-isolated systems, it *can* be the case that the standard statistical mechanical probability distribution does reflect the agent's credences about the microstate of the system, considered as a marginal distribution derived from a distribution about the sys-

[10]See, *e.g.*, Gibbs (1902, pp. 181–83), Khinchin (1949, §20), Thompson (1988, §2.4).

[11]This follows from a theorem of considerable generality due to Popescu et al. (2006), which encompasses restrictions on the global state other than the usual restriction to a fixed energy subspace. Significant earlier work related to this issue includes Goldstein et al. (2006). See Lloyd (2006) for discussion, and further references.

tem and environment. Does this lead to disastrous retrodictions? The canonical distribution, applied at time t_1, would yield high probability that the system was in equilibrium at time t_0, some short time before, if it were further stipulated that the system was isolated in the interim. If it was not isolated, then the dynamics of the system alone entail nothing about its state at earlier times; we need also to specify how its environment acted upon it. Our agent can, without contradiction, have credences about state of affairs at t_1 such that her credences about the system are closely approximated by the canonical distribution (an equilibrium distribution), while having credences about the state of system+environment that yield, when back-evolved, high probability to the system not having been in equilibrium at time t_0. Once again, we can get away with using an equilibrium distribution for the state of the system at t_1, without encountering disastrous retrodictions.

Interventionist approaches to statistical mechanics emphasize the fact that no system is entirely isolated, and hold that it is only when we consider a system interacting with its environment, and restrict attention to the state of the system, that we are justified in using equilibrium probability distributions. It should be noted that interventionist approaches also invoke means-relative considerations. The distinction between a system and is environment is not a matter of fundamental physics. The system is taken to be a system of interest to us, and the rationale for disregarding the degrees of freedom of the environment is that our observations will be restricted to those performed on the system. The chief difference between such an approach and the approach taken in this paper is that, on the present approach, there is a role for equilibrium distributions, even in the idealized case of an isolated system.

There is another way for the entropy of a credence-function to increase. As mentioned, our agent will typically not know the precise dynamics of the system, and, even if known, will typically be unable to perform the calculation. Thus, the agent will be uncertain about what the result is of evolving her credence-distribution ρ_0 about the state of the system at t_0, via the actual dynamics, from t_0 to t_1. She may entertain conjectures about what the result would be, and, provided that her credences satisfy the PPE, her credences about states of affairs at t_1 will be a weighted average of the candidates for what the time-evolute of ρ_0 is. If she is sufficiently uncertain about this, then her credence-distribution about states of affairs at t_1 might be uniform, or nearly so, over the region of phase space accessible to the system at this time. Again, there is no temptation to make disastrous retrodictions. Each of the candidates for the evolute of ρ_0 back-evolves, via the conjectured dynamics associated with it, into ρ_0, even though the mixture of these candidates is invariant under each candidate evolution.

In case this sounds confusing, here's another way of expressing the point. Suppose there is some set $\{\mathcal{T}_\alpha\}$ of dynamical maps, such that Alice believes that one of these yields the actual dynamics of the system. Suppose that her credence-distribution over this set is given by a density $\mu(\alpha)$. Then her credence about the state of affairs at t_1 is given by

$$\bar{\rho}_1 = \int \mathcal{T}_\alpha \rho_0 \, d\mu(\alpha).$$

The averaged distribution $\bar{\rho}_1$ could be an equilibrium distribution. This would mean that, for each $\mathcal{T}_\alpha$, $\bar{\rho}_1$ is invariant under $\mathcal{T}_\alpha$, and hence $\mathcal{T}_\alpha^{-1}$ also leaves it invariant. But, though $\bar{\rho}_1$ is an invariant distribution that represents Alice's credences about the state of the system at t_1, it doesn't follow that it reflects her credences about the state of the system at t_0, as, for each α, $\mathcal{T}_\alpha \rho_0$ back evolves, via $\mathcal{T}_\alpha^{-1}$, into ρ_0.

4 What is Thermodynamics?

Our account of statistical mechanical probabilities combines physical considerations with considerations about feasible measurements and reasonable credences. When the latter are introduced in the context of physical theory, there is resistance from some quarters. We can expect to hear objections that notions such as measurement and credence have no place in a fundamental physical theory.

A goal of statistical mechanics is to recover thermodynamics in a macroscopic limit. In discussing what sorts of consideration are and aren't appropriate to introduce into statistical mechanics, it is worthwhile to pause and consider the nature of thermodynamics.

Thermodynamics begins with a distinction between two modes of energy transfer. Energy may be transferred from one system to another as heat, or via work being done by one system on another. The First Law of Thermodynamics says that any change in total energy of a system can be partitioned into these modes of energy exchange. The Second Law says that heat extracted from a system cannot be converted without residue into work.

If heat were a substance, then the heat content of a system would be a function of the state of a system, and there would be no question of how to partition a change of energy into a part due to heat exchange and a part due to work done. A fundamental fact of thermodynamics is that energy that enters a body as heat can be extracted as work, and *vice versa*. On the kinetic theory of heat, heating and doing work are both processes in which parts of one body act on parts of the other to change their state of motion. The difference is that, when we do work on a body, say, by raising a weight, we move its parts in a systematic way, that we can keep track of; in heating, motion is imparted to the molecules of a body in a haphazard way, and the added energy is quickly distributed among the molecules that make up the body. There is no way to keep track of it in such a way as to wholly recover this energy as work.

This suggests that the distinction between heat and work is relative to the means available to us; it is a matter of what *we* can keep track of. This, in fact, was Maxwell's view. "Available energy is energy which we can direct into any desired channel. Dissipated energy is energy we cannot lay hold of and direct at pleasure, such as the energy of the confused agitation of molecules which we call heat" (Maxwell 1878a, p. 221; Niven 1965, p. 646).

In the opening section of his *Philosophical Essay on Probability*, Laplace famously invited the reader to consider a being that knew the precise state of the world at an instant and all the laws of nature, and was capable of performing the requisite calculations; "for it, nothing would be uncertain and the future, as the past, would be present to its eyes" (Laplace, 1951, p. 4). The point of this passage is to explain why, in spite of the presumed deterministic nature of the laws of physics—which Laplace seems to have taken as an *a priori* truth—there should be such a subject as probability theory. Laplace's answer is that, though such an intelligence would have perfect knowledge of all events, past, present, and future, our own state of knowledge will always remain infinitely removed from this ideal. Hence, we must use probability theory in order to cope, systematically, with less-than-perfect knowledge.

Maxwell, in his *Theory of Heat* (1871, pp. 308–309), invites us to imagine a being that could keep track of and manipulate individual molecules. The actions of such a being would not, according to Maxwell, be subject to the Second Law of Thermodynamics. Moreover, the very concepts need to express the law would not be ones that would occur to it; "we have only to suppose our senses sharpened to such a degree that we could trace the motions of molecules as easily as we now trace those of large bodies, and the distinction between work and heat would vanish, for the communication of heat would be seen to be a communication of energy of the same kind as that which we call work" (Maxwell, 1878b, p. 279). Just as Laplace's demon would have no use for probability theory, Maxwell's demon would have no use for the science of thermodynamics.

Thermodynamics, in its very formulation, requires a distinction between those aspects of a system that are within the scope of our knowledge and control, and those that are not. It is because of our inability to keep track of and manipulate individual molecules that we regard some processes as dissipative. This distinction is reflected in our statistical mechanical treatment of macroscopic systems. In statistical mechanics, we distinguish between variables that we regard as known and use to define a thermodynamic state, and those over which we define a probability distribution. We also distinguish between types of interaction between a system and the rest of the world. Consider a gas in a container that is in thermal contact with a heat bath. The gas exerts a pressure on the walls of the container; this is due to the forces of repulsion between the molecules of the gas and the walls of the container. But the gas is also in thermal equilibrium with the walls of the container, which, if non-insulating, may conduct heat from a heat bath. The walls have finite temperature; the molecules that make up the walls are fluctuating about their equilibrium position, and this means that the forces exerted on a molecule that approaches the wall will also be subject to fluctuations about their mean values. Implicitly,[12] we partition the interaction of the gas with the walls of the container into two terms, a mean value associated with the macroscopic position of the walls, and a term responsible for thermalization of the gas. It is via the former that we do work on the system, by manipulating the macroscopic position of the walls; energy transfer via the latter is regarded as heat transfer. This distinction between two sorts of interaction is essential to any statistical mechanical construal of a distinction between heat and work.

If the goal of statistical mechanics is to recover thermodynamics, and if the very concepts with which the laws of thermodynamics are formulated only make sense with reference

[12] And sometimes explicitly; see, *e.g.*, Thompson (1988, §2.5).

to some agent's capacity to keep track of and manipulate molecules, then it should come as no surprise that such considerations enter into the formulation of statistical mechanics. Yet there are laws of thermodynamics, which talk about an agent can and cannot achieve, and these depend only on very general assumptions about what the agent is able to do—according to Maxwell, the second law of thermodynamics is valid only insofar as molecules are not being manipulated individually or in small numbers. The conception of probability outlined in this paper—incorporating both epistemic and physical considerations—thus seems appropriate for the goal of recovering thermodynamics, so conceived.

5 Explaining Relaxation to Equilibrium

The preceding will not still all qualms about introducing subjective considerations into a physical context. For one thing, we commonly apply statistical mechanics, not only to our laboratory manipulations, but also to situations in which no agents are present, including the early universe. In addition, it looks as if we are giving up on the goal of explaining thermodynamic behaviour of physical systems, as our knowledge and beliefs about such systems are surely not to to be included in an explanation of why they behave as they do![13]

It is certainly correct that considerations of limitations of our knowledge and manipulative prowess are out of place in explanations of the behaviour of systems (except those that we happen to be manipulating). A system behaves as it does because its dynamics, together with initial conditions. Explanations of relaxation to equilibrium will have to involve an argument that the dynamics, together with initial conditions of the right type, yields that behaviour, plus an explanation of why the sorts of physical processes that give rise to the sorts of systems considered don't produce initial conditions of the wrong type (or rather, don't *reliably* produce initial conditions of the wrong type). Nothing that has been said in this paper about probabilities in statistical mechanics touches directly on the problem of explaining thermodynamic behaviour.

However, there is a connection. The processes that are responsible for relaxation to equilibrium are also the processes that are responsible for knowledge about the system's past condition of non-equilibrium becoming useless to the agent. Thus, an explanation of relaxation of equilibrium is likely to provide also an explanation of washing out of credences about the past. Moreover, an explanation of why no process reliably produces initial conditions that lead to anti-thermodynamic behaviour would also explain the reasonableness of credences that attach vanishingly small credence to such conditions. Our judgments about what sorts of processes occur in nature and our judgments about what sorts of credences are reasonable for well-informed agents are closely linked; if there were processes that could reliably prepare systems in states that lead to anti-thermodynamic behaviour, then it would not be unreasonable for an agent to attach non-negligible credence to the system having been prepared in such a state, and we would adjust our judgments about what are and are

[13]For a particularly vivid expression of this point, see Albert (2000, pp. 54–65).

not reasonable credences accordingly.

Furthermore, reflection on the considerations above help us to see what is wanted from such an explanation. If it can be shown that, given fairly weak constraints on an initial probability distribution and on the dynamics, the distribution evolves into one that is, with respect to probabilities of results of macroscopic measurements, effectively indistinguishable from an equilibrium distribution, this would show that weak assumptions about the initial state and the dynamics entail that the system evolves into a state that is macroscopically an equilibrium state. Of particular note in this regard is a series of recent results regarding approach to equilibrium in quantum mechanics. For a very broad class of Hamiltonians (namely, those with nondegenerate energy gaps), the reduced state of a small subsystem of a large quantum system will equilibrate (Linden et al., 2009), provided only that the state of the large system be composed of a large number of energy eigenstates. The equilibrium state is independent of the precise initial state of the bath. A related result due to Goldstein et al. (2010) demonstrates approach to thermal equilibrium of a small subsystem for "typical" Hamiltonians. One expects that there are ways to extend these results to classical mechanics; to work out exactly what assumptions about the classical system are required would take some care but would likely yield insight into conditions of equilibration in classical mechanics.

6 Conclusion

Consideration of epistemic chances resolves the puzzles concerning the status of probabilities in statistical mechanics. It leaves us with a well-motivated research programme, namely, examination of the conditions under which the dynamics of a system will yield almost-objective probabilities—not in the infinite long run, but in finite time. (Results concerning limiting behaviour will, of course be relevant, as they can often be rendered informative about finite time behaviour.)

The familiar dichotomy of subjective and objective probability does not fit well with statistical mechanics, which requires use of probabilistic concepts though the state evolution is deterministic (this remains true of quantum statistical mechanics as currently practiced, which deals with deterministic, unitary evolutions rather than dynamical state reduction), and turns assumptions about probabilities into verifiable empirical predictions. This is not surprising if we consider thermodynamics, the science that it is the goal of statistical mechanics to recover in approximation, to involve a mix of physical considerations and considerations regarding what it is in our power to keep track of and manipulate. Epistemic chances, whose very characterization requires both considerations of epistemic limitations and physical dynamics, seem to be just what are required for the purpose.

7 Acknowledgments

References

Albert, D. (2000). *Time and Chance*. Cambridge: Harvard University Press.

Berkovitz, J., R. Frigg, and F. Kronz (2006). The ergodic hierarchy, randomness and Hamiltonian chaos. *Studies in History and Philosophy of Modern Physics 37*, 661–691.

Gibbs, J. W. (1902). *Elementary Principles in Statistical Mechanics: Developed with Especial Reference to the Rational Foundation of Thermodynamics*. New York: Charles Scribner's Sons.

Goldstein, S., J. L. Lebowitz, C. Mastrodonati, R. Tumulka, and N. Zanghì (2010). Approach to thermal equilibrium of macroscopic quantum systems. *Physical Review E 81*, 011109.

Goldstein, S., J. L. Lebowitz, R. Tumulka, and N. Zanghì (2006). Canonical typicality. *Physical Review Letters 96*, 050403.

Hacking, I. (1971). Equipossibility theories of probability. *British Journal for the Philosophy of Science 22*, 339–355.

Hacking, I. (1975). *The Emergence of Probability*. Cambridge: Cambridge University Press.

Jackson, E. A. (1968). *Equilibrium Statistical Mechanics*. New York: Dover Publication, Inc.

Khinchin, A. I. (1949). *Mathematical Foundations of Statistical Mechanics*. New York: Dover Publications.

Laplace, P.-S. (1814). *Essai Philosophique sur les Probabilités*. Paris: Courcier. English translation in Laplace (1951).

Laplace, P.-S. ([1902] 1951). *A Philosophical Essay on Probabilities*. New York: Dover Publications. Translation of Laplace (1814).

Lewis, D. (1980). A subjectivist's guide to objective chance. In R. C. Jeffrey (Ed.), *Studies in Inductive Logic and Probability*, Volume II, pp. 263–93. University of Caifornia Press.

Linden, N., S. Popescu, A. J. Short, and A. Winter (2009). Quantum mechanical evolution towards thermal equilibrium. *Physical Review E 79*, 061103.

Lloyd, S. (2006). Excuse our ignorance. *Nature Physics 2*, 727–728.

Machina, M. J. (2004). Almost-objective uncertainty. *Economic Theory 24*, 1–54.

Maxwell, J. C. (1871). *Theory of Heat*. London: Longmans, Green, and Co.

Maxwell, J. C. (1878a). Diffusion. In *Encyclopedia Britannica* (Ninth ed.), Volume 7, pp. 214–221. Reprinted in Niven (1965, pp.625–646).

Maxwell, J. C. (1878b). Tait's "Thermodynamics". *Nature 17*, 257–259, 278–280. Reprinted in Niven (1965, pp. 660-671).

Niven, W. D. (Ed.) (1965). *The Scientific Papers of James Clerk Maxwell*, Volume Two. New York: Dover Publications. Reprint of Cambridge University Press edition of 1890.

Plastino, A. and A. Daffertshofer (2008). Classical typicality of the canonical distribution. *Europhysics Letters 84*, 30006.

Poincaré, H. (1912). *Calcul des probabilités.* Paris: Gauthier-Villars.

Popescu, S., A. J. Short, and A. Winter (2006). Entanglement and the foundations of statistical mechanics. *Nature Physics 2*, 754–758.

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

Thompson, C. J. (1988). *Classical Equilibrium Statistical Mechanics.* Oxford: Clarendon Press.

Tolman, R. C. (1938). *The Principles of Statistical Mechanics.* Oxford: Clarendon Press. Reprint, Dover Publications, 1979.

von Plato, J. (1983). The method of arbitrary functions. *The British Journal for the Philosophy of Science 34*, 37–47.