

# The Justification of Probability Measures in Statistical Mechanics\*

Kevin Davey<sup>†</sup>

---

According to a standard view of the second law of thermodynamics, our belief in the second law can be justified by pointing out that low-entropy macrostates are less probable than high-entropy macrostates, and then noting that a system in an improbable state will tend to evolve toward a more probable state. I would like to argue that this justification of the second law is unhelpful at best and wrong at worst, and will argue that certain puzzles sometimes associated with the second law are merely artifacts of this questionable justification.

---

**1. The Standard Story.** If we squirt some colored ink into a closed container of moving water, we are justified in expecting that the ink will eventually disperse itself evenly through the water. Indeed, if we prepare any system in a state of low entropy, and then isolate the system, we are justified in expecting that the system will experience an increase in entropy. In virtue of what are these beliefs justified? Such beliefs can, of course, be justified on purely inductive grounds—systems like this have behaved similarly in the past. But one might also think that such beliefs ought to be able to be justified on the basis of more fundamental physical and mathematical principles. In this paper, I would like to discuss and criticize a common justification of this latter sort.

Inasmuch as it is possible, I will use terms like ‘justification’ only in their most uncontroversial sense. When I say that a belief is justified, all that I mean is that we have adequate grounds for believing it. I take no side, for instance, on the question of whether those grounds are to be spelled out in internalistic or externalistic terms.<sup>1</sup> It will suffice to use

\*Received October 2006; revised October 2007.

<sup>†</sup>To contact the author, please write to: Department of Philosophy, University of Chicago, 1115 E. 58th St., Chicago, IL 60637; e-mail: kjdavey@uchicago.edu.

1. Some modest restrictions, however, must be placed on externalist theories of justification in order for my argument to work. This is discussed in Section 3.

Philosophy of Science 75 (January 2008) pp. 28–44. 0031-8248/2008/7501-0002\$10.00  
Copyright 2008 by the Philosophy of Science Association. All rights reserved.

facts and intuitions about justification that I think are common to the internalist, externalist, and others alike.

The challenge implicitly posed in the first paragraph is to explain why we are justified in believing the second law of thermodynamics, given only more fundamental physical and mathematical principles. For our purposes, the second law of thermodynamics states that the entropy of an isolated system will increase until the system reaches its equilibrium state, which is a state of maximum entropy. Once in this equilibrium state, the system will remain there (or at least will remain there for a very long time.)

A standard (though schematic) argument for the second law is:

**Standard Story.** Low-entropy macrostates occupy a tiny portion of phase space, and so it is extremely improbable that a system will find itself in a low-entropy macrostate. High-entropy macrostates occupy essentially all of phase space, and so it is extremely probable that a system will find itself in a high-entropy macrostate. If a system has an extremely probable and an extremely improbable state, then whenever the system is in the extremely improbable state, it will very likely find itself in the extremely probable state shortly thereafter, and will likely remain in the extremely probable state for some time. (That, after all, is just what it *means* to call a state extremely probable.) Thus, a system in a low-entropy macrostate is overwhelmingly likely to evolve into a high-entropy macrostate, and remain in that high-entropy macrostate for some time.

The primary goal of this paper will be to argue that this story is inadequate at best and wrong at worst as an explanation of why we are justified in expecting isolated systems to obey the second law.

In Section 3, I argue that the Standard Story rests critically on what I will call a *Probability Principle*. A Probability Principle is any claim of the form:

**Probability Principle.** If we know that a system is in macrostate  $M$ , then we are justified in describing the microstate of the system with the probability measure  $\mu_M$ , where  $\mu_M$  is defined in terms of  $M$  as  
 . . .

Different procedures for defining  $\mu_M$  correspond to different Probability Principles.

In Section 3, I argue that no Probability Principle can be correct, and that the Standard Story is thus fundamentally flawed. In Section 4, I argue that this sheds light on some problems traditionally associated with statistical mechanics. For instance, one problem, sometimes grouped together with reversibility objections, is that statistical mechanics tells us

that an isolated system in a medium-entropy state is much more likely to have had a high-entropy past than a low-entropy past, contrary to our experience. Some sort of maneuvering is therefore needed in order to salvage our ability to reason about the past using statistical mechanics. Precisely what sort of maneuvers are needed is a subject of some dispute. Even once such maneuvers are made, however, many accounts leave us having to say that the universe began its life in an exceedingly improbable state. We are then left wondering whether and how we are to explain this initial state. My arguments, however, will show that there is no clear sense in which statistical mechanics gets the past wrong, and no clear sense in which we are forced to say that the universe began its life in an exceedingly improbable state. Indeed, there is no clear sense in which low-entropy states are a priori improbable, and high-entropy states a priori probable. Many associated philosophical problems then disappear—though of course, many remain. In Section 5, as an ‘appendix’ of sorts, I will discuss a few reasons why some have felt tempted to adopt Probability Principles, and will explain why these reasons are misguided.

A caveat is in order before beginning. It is not clear that the Gibbsian will justify belief in the second law by giving the Standard Story. This is so for several reasons. First of all, for many Gibbsians, what makes a system in a low (coarse grained) entropy state tend towards a high (coarse grained) entropy state is not the fact that improbable states are likely to be followed by probable states, but rather facts about the dynamics of mixing. Second of all, the Gibbs entropy of a system is a functional of the probability distribution appropriate to the system, rather than being directly defined from the macrostate itself. As such, it is not clear that a claim such as ‘low-entropy macrostates are improbable because they occupy a tiny portion of phase space’ even makes sense to a Gibbsian. I do think, however, that the Standard Story is implicit in many so-called Boltzmannian accounts of the second law (e.g., Albert 2000), and so this paper is restricted to a discussion of such accounts.

**2. Justifying the Standard Story.** Crucial to the standard story is the following:

**Basic Claim.** If a region of phase space is large, then it is probable for a physical system to find itself there.

Why should we accept this principle?

There seems to be no good reason to think that the Basic Claim must be true of dynamical systems in general. Say that a subset  $S$  of phase space is *invariant* if time-evolution can neither take any point from outside  $S$  into  $S$ , nor any point inside  $S$  outside of  $S$ . Suppose that phase space can be divided into at least two disjoint invariant sets  $S_1$  and  $S_2$ . If the

present microstate of a physical system lies in  $S_2$ , then it will never find itself in  $S_1$ , no matter how large  $S_1$  is. In what sense then is region  $S_1$  *probable*? One suggestion might be that if  $S_1$  is large, very few systems will lie in  $S_2$ , and in this sense,  $S_1$  is probable. But how is *this* claim to be justified without appealing to the Basic Claim? We have led ourselves in a circle. Moreover, even if we simply take at face value the claim that it is improbable for a system to find itself outside of  $S_1$ , trouble looms. For suppose some macrostate  $M$  lies entirely outside  $S_1$ , and that all the maximum-entropy microstates lie in  $S_1$ . Then, with certainty, any system in macrostate  $M$  will never evolve into a maximum-entropy state, contrary to the second law of thermodynamics.

Faced with these and other problems, most people acknowledge that any defense of the Basic Claim requires us to restrict our attention to certain types of dynamical systems. Generally, people suppose that the system must be *ergodic*; that is, that phase space cannot be separated into two invariant subspaces of nonzero measure.<sup>2</sup> Armed with this assumption, it turns out that certain time averages can be equated with space averages, and the usual development of equilibrium statistical mechanics proceeds swimmingly.

It is possible to make all sorts of assumptions related to (but distinct from) ergodicity, and end up in the same happy situation. For our purposes, we will not work with ergodicity, but will work with an even stronger assumption, namely, the mixing hypothesis:

**Mixing Hypothesis.** Let  $S$  be phase space, and let  $\mu$  be the standard Lebesgue measure on  $S$ , normalized so that  $\mu(S) = 1$ . Let  $M$  and  $Q$  be measurable subsets of  $S$ , and let  $M_t$  be the  $t$ -second evolution of  $M$ . Then

$$\lim_{t \rightarrow \infty} \mu(M_t \cap Q) = \mu(M)\mu(Q). \quad (1)$$

To see how the Mixing Hypothesis gives us the second law, let  $Q$  be the set of equilibrium states of the system (so that  $\mu(Q) \approx 1$ ), and let  $M$  be an arbitrary macrostate (assumed to have nonzero measure). Then

$$\lim_{t \rightarrow \infty} \frac{\mu(M_t \cap Q)}{\mu(M)} \approx 1. \quad (2)$$

This is usually interpreted as saying that almost all microstates in  $M$  eventually end up in  $Q$ . It follows that a system in macrostate  $M$  is almost certain to end up in  $Q$ . Thus, we have established a version of the second law. Moreover, insofar as it is the largeness of  $Q$  that guarantees that our

2. For an expository account, see Sklar 1993, Chapter 5, Sections I and II.

system will almost certainly end up in  $Q$ , this argument also vindicates the Basic Claim and thus the Standard Story.

The Mixing Hypothesis (and its close relatives) have been criticized by Sklar (1993), Earman and Redei (1996), and others. These authors suggest that there is no evidence that such assumptions hold of realistic physical systems, and that there is perhaps even evidence that such assumptions fail. I take these criticisms to be reasonable.

I think, however, that there is a more basic error that occurs in this justification of the second law. Suppose we have a system for which the Mixing Hypothesis holds. Do we then really have the materials with which we can justify the Basic Claim, and thus the Standard Story? I claim not.

To see the main problem, we ask why we should interpret (2) as saying that a system in macrostate  $M$  is extremely likely to end up in  $Q$  at some point in its future. The argument for this interpretation of (2) is as follows: suppose our system is in macrostate  $M$  at  $t = 0$ . Choose a probability measure  $\mu_M$  over the set of all possible microstates that is uniform over  $M$ , and which vanishes elsewhere. That is to say, if  $T \subset S$ , then the probability that our system is in a microstate in  $T$  is given by

$$\mu_M(T) = \mu(M \cap T)/\mu(M), \quad (3)$$

where  $\mu$  is the standard Lebesgue measure. Given any  $T \subset S$ , and any real  $r$ , let  $T_r$  be the set of microstates which,  $r$  seconds ago, were elements of  $T$ . Given this notation and this choice of probability measure, the probability that  $t$  seconds from now our system will lie in  $Q$  is

$$\mu_M(Q_{-t}) = \mu(M \cap Q_{-t})/\mu(M). \quad (4)$$

Using Liouville's Theorem,  $\mu(M \cap Q_{-t}) = \mu(M_t \cap Q)$ . Thus, the probability that  $t$  seconds from now our system will lie in  $Q$  is

$$\mu(M_t \cap Q)/\mu(M). \quad (5)$$

The Mixing Hypothesis tells us that this quantity approaches a number close to 1 as  $t$  increases. Thus, our system is overwhelmingly likely to end up in  $Q$  at some point in its future.

But the cogency of this argument depends on our choice of the measure  $\mu_M$  as the relevant probability measure with which to describe the present state of the system. In particular, the following principle is being invoked:

**Probability Principle 1.** If we know that a system is in a macrostate  $M$ , then we are justified in describing the present state of the system with the probability measure  $\mu_M$  that is uniform over  $M$ , but which vanishes elsewhere.

Without an assumption like Probability Principle 1, we are in no position

to interpret (2) as telling us that a system in the macrostate  $M$  is extremely likely to end up in  $Q$  in the future. To see why, let  $M$  be a macrostate, and let  $x \in M$  be a microstate which never passes through the equilibrium macrostate  $Q$ .<sup>3</sup> Define  $\mu_x$  to be a probability measure on  $M$  which is supported only on  $x$ ; that is, for any subset  $T$  of phase space,

$$\mu_x(T) = \begin{cases} 1 & \text{if } x \in T \\ 0 & \text{if } x \notin T \end{cases}. \quad (6)$$

Using this measure of probability, a system in the macrostate  $M$  is *certain* to remain outside of  $Q$  forever, even though (2) holds. Thus, it is only if we draw some sort of connection between Lebesgue measure and the actual probability measure pertinent to the present state of the system that we can take (2) to be telling us that a system in the macrostate  $M$  is extremely likely to end up in  $Q$ . This sort of connection is explicitly drawn in Probability Principle 1.

So even if the Mixing Hypothesis holds, we get to tell something like the Standard Story only if something like Probability Principle 1 is assumed. But is Probability Principle 1 correct? The answer is no. Let  $M$  be a non-equilibrium macrostate of a system, and suppose we know that our system is in  $M$ . Almost all microstates in  $M$  have a higher-entropy past. If we are justified in describing the microstate of the system with the probability measure  $\mu_M$  described by Probability Principle 1, then we are justified in believing that our system almost certainly had higher entropy in the recent past. But we have strong inductive grounds for believing that our system did not have higher entropy in the recent past. Therefore, Principle 1 is false. (Essentially this line of argument is developed in Albert 2000, Chapter 4.)

The natural thing to do at this point is modify Probability Principle 1. Here is a proposal that circumvents the problem:

**Probability Principle 2.** If we know a system is in a macrostate  $M$ , then we are justified in describing the microstate of the system with the probability measure  $\mu_{M^*}$  that is uniform over  $M^*$ , but which vanishes elsewhere, where  $M^*$  is the set of microstates  $x \in M$  such that  $x$ 's entropy was lower in the recent past.

(Here, by the 'recent past', we mean the interval of time  $[-s, 0]$  over which we have good reason for believing that our system has been isolated.) Here is another proposal:

3. For instance, in the case of hard spheres scattering in a closed rectangular container (one of the few cases for which the Mixing Hypothesis may be proved to hold), let  $x$  be a microstate in which all particles have velocities parallel to a fixed wall of the container, and in which no collisions between particles ever occur.

**Probability Principle 3.** Same as Probability Principle 2, except  $M^*$  is now the set of microstates  $x \in M$  compatible with the early universe being in some particular very low-entropy macrostate at some particular time.

And here is another:

**Probability Principle 4.** Same as Probability Principle 2, except  $M^*$  is now the set of microstates  $x \in M$  that behave in a way that is compatible with all our inductively justified expectations about the past and future of the system.

Regardless of which of these principles is chosen, we can argue from the Mixing Hypothesis that

$$\lim_{t \rightarrow \infty} \frac{\mu(M_t^* \cap Q)}{\mu(M^*)} \approx 1. \quad (7)$$

Using the relevant Probability Principle, it follows that a system in macrostate  $M$  is very likely to end up in  $Q$ , in virtue of  $Q$ 's size, and so the Basic Claim and Standard Story are vindicated.

But are any of Probability Principles 2, 3, or 4 true? I claim that they are not. In fact, I claim that all principles of the form

**Probability Principle.** If we know a system is in macrostate  $M$ , then we are justified in describing the microstate of the system with the probability measure  $\mu_M$ , where  $\mu_M$  is defined in terms of  $M$  as . . .

are false. But without a Probability Principle of this sort, it is difficult to see how the knowledge that a system obeys the Mixing Hypothesis can be relevant in vindicating the Basic Claim or the Standard Story. I move to my argument for the falsehood of Probability Principles now.

**3. Against Probability Principles.** The main problem with any Probability Principle is that, according to such a principle, knowing the current macrostate  $M$  of a system (and having some background knowledge about the universe and the way in which physical systems generally behave) is sufficient for us to have a justified belief about the probability with which our system lies in any subregion of phase space. But it is perfectly possible for us to know the current macrostate of a system, without being in a position to form a justified belief about the probability measure with which to describe the microstate of the system. Or to say the same thing differently—it is possible for us to know the current macrostate of a system, without having any grounds for believing that any particular probability measure is in fact the correct probability measure with which to describe the microstate of the system. Therefore, all Probability Principles are false.

Consider the following situation. Suppose a scientist brings in a glass box, which we observe to be filled with a gas in its equilibrium macrostate  $M$ . Suppose the scientist tells us that 10 minutes ago the gas was in equilibrium, and at that time he made a choice between:

**Option A.** Leaving the gas undisturbed, in its equilibrium macrostate, or

**Option B.** Forcing the gas out of its equilibrium macrostate, into a specific macrostate that typically takes 10 minutes to relax back into equilibrium.

The scientist will not tell us which option he chose, and knowing that the system is presently in its equilibrium macrostate does not help us deduce which choice was made. The scientist does tell us, however, that 1 hour ago he selected a real number  $\alpha$  with  $0 < \alpha < 1$  (by some unknown procedure) and that he then devised some random process that guaranteed that he chose Option A with probability  $\alpha$ , and Option B with probability  $1 - \alpha$ .

For no value of  $\alpha$  are we justified in believing that  $\alpha$  was the probability chosen by the scientist. We may assume that we know nothing about the scientist, and nothing about the procedure by which he chose  $\alpha$ . (We do not even know whether his selection of  $\alpha$  was truly random, so we are not even entitled to claim a uniform probability distribution for  $\alpha$ .<sup>4</sup>) Again, it does not matter whether we are working with an internalistic or externalistic conception of justification, as we can imagine that we are in a situation in which we are not even justified by any vaguely plausible externalistic standards in believing that some particular  $\alpha$  is in fact the  $\alpha$  chosen by the scientist.<sup>5</sup>

Suppose, however, that some Probability Principle tells us that we are justified in describing the present state of our system with a probability measure  $\mu^*$ . Let  $N$  be the set of microstates that, 10 minutes ago, were in the particular out of equilibrium macrostate referred to in Option B. From the fact that we are justified in using the probability measure  $\mu^*$ , it follows that we are justified in believing that  $1 - \alpha = \mu^*(N)$ . But we are not justified in believing that  $1 - \alpha = \mu^*(N)$ , because *no* belief about the exact value of  $\alpha$  is justified. Therefore, the Probability Principle is

4. I assume that we are not justified in using the Principle of Indifference to posit a uniform probability distribution for  $\alpha$ . For some critical comments on the Principle of Indifference, see Section 5.1.

5. Obviously, an externalism according to which truth is enough for justified belief will undercut my argument here, but an externalism of this sort is not plausible.

false. So no Probability Principle can be correct—and it is worth noting that this is so even when the macrostate  $M$  is an equilibrium macrostate.<sup>6</sup>

One might try to argue that because the second law talks about isolated systems, the only notion of probability that matters is one that describes systems while they are isolated. But even the probable behavior of isolated systems will depend on the way in which the systems are prepared and how long they are allowed to stay isolated. As the previous examples show, such contingencies can make it the case that even low-entropy states turn out to be probable.

The lesson here is that the correct probability measure with which to describe the state of a system is, in general, determined by all sorts of facts about the past of the system, including, for instance, the way in which the system was prepared. Such facts about the past of the system need not be known by us, nor need we be in a position to form justified beliefs about such facts. Because of this, we will not generally be justified in making claims about the correct probability measure with which to describe the present state of the system.<sup>7</sup>

Of course, this is not to deny that there might be cases in which we *can* have a justified belief about which probability measure describes the current state of a system;<sup>8</sup> nor it is to deny that we can always make an *unjustified* claim or conjecture about the probability measure describing the current state of the system. All I intend to deny is that there is some sort of procedure that can be used to move from knowledge of the present macrostate of a system to a justified probability measure describing the state of the system.

It is worth noting that Albert's (2000) main proposal also falls prey to the sort of argument given above, even though it is not clear that he is advancing a Probability Principle as I have defined it. Albert suggests that if a system is currently in a macrostate  $M$ , we ought to describe the

6. One can easily generate similar examples in which the final macrostate  $M$  is not an equilibrium macrostate. For instance, perhaps the scientist presents us with a glass of water with a thin film of ice on the surface, and tells us that 1 hour ago, he either placed six ice cubes in a glass of water at room temperature (with probability  $\alpha$ ) or seven slightly less cold ice cubes in a glass of water at room temperature (with probability  $1 - \alpha$ ).

7. This phenomenon also arises in the case of non-Markovian dynamics. Here, the future behavior of a system depends not only on the present state of the system, but also on the history of the system. If we know the present macrostate of such a system, but do not know its history, we will generally not be in a position to specify a probability measure over all possible present microstates.

8. For instance, we might have strong inductive grounds for describing the state of a coin tossed in the air with a probability measure that is uniform over the set of all possible angular orientations.

possible microstates of the system with a probability distribution that is uniform over the subset of  $M$  compatible with the universe having started in the particular low-entropy state it started in, and also “compatible with whatever other information—either in the form of laws or in the form of contingent facts—we happen to have” (Albert 2000, 96). If Albert is right, then there is some particular probability distribution I ought to assign to the box of gas the scientist puts before me. It follows that I am justified in believing that  $\alpha$  has some particular value. Because I am not justified in ascribing any specific value to  $\alpha$ , Albert’s proposal cannot be correct.<sup>9</sup>

One might think that the problem here is that we need to adopt a principle that includes information about the past of the system. For instance, consider:

**Historic Probability Principle.** Suppose we know that a system has been isolated during the interval of time from  $t = -r$  to  $t = 0$ . Suppose also that at any time  $s$  in this range, the system is known to have been in macrostate  $M_s$ . Then we are justified in describing the microstate of the system with the probability measure  $\mu$ , where  $\mu$  is defined in terms of  $\{M_t | t \in [-r, 0]\}$  as follows . . .

The problem, however, is that in general we do not know the past macroscopic history of systems of interest, and so relying on a Historic Probability Principle severely limits the sorts of situations in which we are justified in expecting a system to obey the second law. But putting this to the side, it is not clear whether (and in what circumstances) knowledge of the complete macroscopic history of an isolated system really puts us in a situation in which we are justified in describing the present microstate of a system with a particular probability distribution. That is an issue for further study.

Finally, one might think that we have to limit somehow the ways in which probability principles are applied, perhaps even in some time asymmetric manner. For instance, Winsberg (2004b) argues that we should apply probability principles to a system only at a certain moment of the system’s history—specifically, at the moment at which the system becomes

9. One might be worried that we have taken Albert out of context here, for on p. 96 he seems to be talking about the probability distribution that describes the current state of the universe *as a whole*. But if Albert’s proposal really applies only to the universe as a whole, then he has at best given us a reason to think that the universe as a whole obeys the second law, and not that isolated subsystems of the universe obey the second law. Given that Albert’s examples earlier in the chapter revolve around small isolated subsystems of the universe, it is reasonable to ascribe to him the sort of view in the present paragraph. For a related critical discussion of Albert on this point, see Winsberg 2004a.

energetically isolated. It is difficult to see how this sort of restriction can be independently justified, as we must surely demand. It is also difficult to see how to apply Winsberg's rule when we do not have justified belief about when exactly energetic isolation occurred and what the macrostate was at that moment, as is the case in the main example of this section. Whether these problems can be addressed is a question I am happy to leave open. For now, I simply assume that there is no compelling reason to adopt a probability principle at one time and not another.

#### 4. Consequences.

*4.1. Entropy and Probability.* What effect does all this have on our understanding of the second law, and in particular, the Standard Story? According to the Standard Story, because low-entropy macrostates occupy a tiny portion of phase space, they should be assigned low probability. But it is difficult to see how this claim can be justified, other than with a Probability Principle. Because Probability Principles are incorrect, the Standard Story falls into jeopardy.

In fact, once we free ourselves from incorrect Probability Principles, the situation for the Standard Story get even worse. I shall argue that the low-entropy states actually found in nature need not be improbable at all. If low-entropy states are not necessarily improbable, then it cannot be correct to say that the second law of thermodynamics is a result of the fact that systems tend to move from improbable to probable states. The Basic Claim underlying the Standard Story is therefore false.

How might a low-entropy state turn out *not* to be improbable? Before we address this question, let us first consider the following counterargument:

We *should* think of low-entropy states as improbable, because it is exceedingly unlikely that an isolated system, in a high-entropy state, will fluctuate out of that state and into a lower-entropy state shortly thereafter.

There is something irrelevant about this argument. Although it is possible for a system in equilibrium to fluctuate out of equilibrium and into a lower-entropy state, very few systems we actually find in low-entropy states are in such states as a result of such a fluctuation. Most of the world's glasses of ice water, for instance, are the products of very deliberate interactions in which ice and water end up mixed, and are not the result of random thermal fluctuations that began with an isolated glass of water. In what sense, then, is the glass of ice water that we actually find in nature in an improbable state? There is no clear sense in which this is so, as the following examples help to make clear.

Suppose we live in a world in which, by government fiat, all good citizens must do everything they can to keep all glasses of water half full of ice, and that the citizenry is largely successful at this. Suppose we find a glass of water with no ice in it. We will be justified in believing that one hour ago, the glass of water probably had ice in it, because experience suggests that most glasses of water almost always have ice in them. The state of the glass of water in which it contains ice is thus highly probable, even though it has low entropy. In this world, when a glass of ice water melts, it moves from a probable state to an improbable state.

By contrast, consider a world in which, by government fiat, all good citizens must do everything they can to keep all glasses free of ice, and that the citizenry is largely successful at this. Suppose we find a glass of water with no ice in it. We will be justified in believing that one hour ago, the glass of water probably had no ice in it, because most glasses of water almost never have ice in them. The state of the glass of water in which it contains ice is thus highly improbable. In this case, when a glass of ice water melts, it moves from an improbable state to a probable state.

One might try to argue that because the second law talks about isolated systems, the only notion of probability that matters is one that describes systems while they are isolated. But even the probable behavior of isolated systems will depend on the way in which the systems are prepared and how long they are allowed to stay isolated. As the previous examples show, such contingencies can make it the case that even low-entropy states turn out to be probable.

The lesson here is that whether a low-entropy state of a system has low probability, and whether a high-entropy state of a system has high probability, depends on the environment in which the system exists, and the sorts of interactions the system is likely to undergo. There is no straightforward, a priori connection between low entropy and low probability. It cannot then be the case that the second law of thermodynamics is just an expression of the fact that an isolated system will tend to move from an improbable to a more probable state.

To all this, a critic might reply that the universe is destined to undergo heat death, and so in the long run, low-entropy states end up being very rare, and hence improbable. But this point is surely irrelevant. If we want to justify inductively statements about the way that subsystems of the universe will behave when the universe is in a particular state far from equilibrium, then we will focus on our experience of the universe while it is in that particular state. The fact that, on cosmic time scales, that particular state far from equilibrium occupies the mere blink of an eye is irrelevant to the justification of statements about what happens when the eye is in fact blinking.

A corollary should be noted. There has been some literature (see, for

instance, Callendar 2004) on the question of whether an especially improbable initial state of the universe is the sort of thing that requires explanation. This discussion is motivated by the fact that the universe is thought to have started in an extremely low-entropy state. If we take extremely low-entropy states to be extremely improbable, then we are forced to say that the universe started in an extremely improbable state, and we must decide how to react to this. But once we abandon the idea that low entropy means low probability, we are no longer forced to say that the universe began in an extremely improbable state. Indeed, once we abandon all Probability Principles, we are no longer even required to take probabilistic claims about the universe's initial state to be *meaningful* (though we are also free to take some such claims to be meaningful and false). This is not to deny that such claims could be meaningful, nor is it to deny that the universe may turn out to have originated in an extremely improbable state. But neither the meaningfulness nor the truth of the claim that the universe originated in an improbable state follow from the mere fact that the universe started in an extremely low-entropy state. The burden rests on someone who nevertheless believes that the universe started in an extremely improbable state to demonstrate otherwise.

4.2. *Does Statistical Mechanics Get the Past Wrong?* Another conceptual problem with statistical mechanics is that it can appear to give us information about the past at odds with our experience. Suppose we adopt Probability Principle 1. We are then forced to say that a system in a medium-entropy state most likely had higher entropy in the recent past. However, we surely have strong inductive grounds for thinking that the system had lower entropy in the recent past, and so we have a problem.

The problem, of course, is with our choice of  $\mu$ . Suppose that a probability distribution,  $\mu$ , that we use to describe the present state of a system entails some proposition  $X$  about the past of our system, where  $X$  is at odds with our experience. The correct conclusion to draw is that we are not justified in using  $\mu$  to describe the present state of the system—for if we were, we would then be justified in believing  $X$ , which we are not.

In general, any choice of probability measure to describe the present state of a system must be justifiable on inductive grounds. So, for instance, if we are justified in describing the present state of a system with a probability measure  $\mu$  for which  $\mu(X) \approx 1$  (for some proposition  $X$ ), then it must be the case that we have strong inductive grounds for believing that  $X$  is true. The only case in which statistical mechanics can tell us to believe something about the past that is inconsistent with our best inductively formed expectations is when we use an unjustified probability measure  $\mu$ .

Only in this relatively uninteresting sense is statistical mechanics capable of getting the past wrong.<sup>10</sup>

It will be instructive to apply these considerations to the ‘skeptical catastrophe’ discussed by Albert (2000, Chapter 6).<sup>11</sup> The problem is this: suppose I have knowledge only of the present macrostate of a small portion of the universe, and that I am presented with a very old photograph of my grandmother. Experience suggests that this photograph was probably formed in a less ragged state (i.e., in a lower-entropy state) some time ago, as the result of an interaction of a camera with my grandmother. If, however, we describe the state of a closed subsystem of the universe with a probability measure that is uniform over its present macrostate, then it turns out to be exceedingly unlikely that the photograph had such a past. Instead, it is much more likely that the photograph is the result of some sort of fluctuation from equilibrium of a piece of photographic paper. There is nothing special about photographs here—essentially the same point can be made about *all* records of the past, including our memories. And so we must conclude that pretty much all our beliefs about the past, insofar as they rest on records, memories and relics, are extremely unlikely to be true. This is a skeptical catastrophe.

Albert resolves this catastrophe by adopting his ‘Past Hypothesis’, but a much simpler solution is available. On one hand, the Probability Principle on which the argument is based is false. The skeptic, however, may reply that once we place ourselves in the position in which we have knowledge only of the present, we then have *no grounds* for thinking that the relevant Probability Principle is false. Ignoring the question of whether there is an unfair shift in the burden of proof here, this reply quickly gets the skeptic in even more trouble. For suppose we have knowledge only of the present macrostate of some chunk of the universe. A probability measure can be justified only on inductive grounds, and so if all our knowledge is knowledge of the present, then we have no good inductive grounds with which to justify *any* probability measure. But without a justified probability measure, statistical mechanics yields no *justified* beliefs about the past. And if statistical mechanics yields no justified beliefs about the past, then statistical mechanics has no consequences capable of conflicting with common sense. Thus, there is no skeptical catastrophe.

**5. Defenses of Probability Principles.** Before finishing, it will be of interest to consider some arguments that might be given for various Probability Principles. We will consider arguments based on the Principle of Indif-

10. Although Albert (2000) appears to draw a similar conclusion, he then goes on to adopt a Probability Principle of the sort criticized earlier.

11. The phrase ‘skeptical catastrophe’ is taken from Albert 2000, 116.

ference, arguments based on the stipulation that a probability measure be stationary, and arguments based on ergodicity or frequentist considerations. Although some of these arguments (such as that based on stationarity) occur mainly in the context of Gibbsian statistical mechanics—which we are not discussing—it is tempting to think that such arguments can also be used to justify Probability Principles. It is in this regard that we examine these arguments.

*5.1. The Principle of Indifference.* What sorts of arguments might be given to support Probability Principle 1? Perhaps the most straightforward way to justify this choice of probability distribution is to invoke the Principle of Indifference, according to which we should choose a uniform probability measure over a set of possibilities whenever we lack a reason to choose a nonuniform probability measure over that set of possibilities. This sort of reasoning can also be used to justify Probability Principles 2, 3 and 4, and even certain Historical Probability Principles.

The Principle of Indifference as thus stated has been much criticized (see especially van Fraassen 1990, Chapter 12).<sup>12</sup> Some critics have pointed out that the Principle of Indifference fails to deliver a unique probability distribution; this has become known as ‘Bertrand’s Paradox’, and is a serious problem for the Principle. An even more compelling, but separate, concern is that the absence of a reason for choosing a nonuniform probability measure surely does not *justify* the choice of *any* probability measure, let alone a uniform probability measure. A similar point has also been made by Sklar (1993, 118–120).

*5.2. Stationarity Arguments.* Let us focus on the case in which our system is in equilibrium. Let  $\mu$  be the probability measure which is uniform over all phase space. Then  $\mu$  has the property that it is ‘stationary’; that is to say, it remains invariant under evolution in time. Furthermore, if the system is ergodic, then  $\mu$  is the *only* stationary probability distribution that assigns probability 0 to sets of measure 0.<sup>13</sup>

Given that we expect the properties of a system in equilibrium to remain constant over time, and given the reasonableness of assigning probability 0 to sets of measure 0, we conclude that if an ergodic system is in equilibrium, then we are justified in using  $\mu$  to describe the possible microstates of the system.

This probability principle is quite modest, for it is only intended to apply to an ergodic system in equilibrium. Nevertheless, it is false. While

12. For a defense of the Principle of Indifference, see Castell 1998.

13. For a discussion of this point, see Sklar 1993, 159–161.

we have strong inductive grounds for thinking that the properties of an isolated system in equilibrium will not change in the future, we have no grounds for thinking that such a system's properties did not change in the past. Many isolated systems presently in equilibrium were isolated in a state out of equilibrium. If we know that an isolated system was prepared in a state out of equilibrium, or if we are not in a position to judge whether a system was out of equilibrium when first isolated, then we are not justified in employing a stationary probability measure to describe the present microstate of the system.<sup>14</sup>

*5.3. Ergodicity and Frequentism.* Finally, suppose a system is ergodic. Fix a subregion  $R \subset S$  of phase space  $S$ , and let  $\alpha = \mu(R)/\mu(S)$ . Then, for almost all points  $x$  of phase space, the proportion of time that a system beginning in microstate  $x$  will spend in  $R$  is just  $\alpha$ . If we identify the probability of a system being in  $R$  with the proportion of time it spends in  $R$ , as a frequency interpretation of probability might suggest, then we must conclude that the probability of our system being found in  $R$  is  $\alpha$ . Thus, subregions of phase space with equal volume must be assigned equal probability. And this means that, given no additional information about the system, we ought to describe the present microstate of the system with a probability distribution that is uniform over all of phase space.

But similar problems arise. Typically, we do not know nothing about a system; the additional knowledge we have about a system will generally allow us to argue that the system is more likely to be in one subregion of phase space than another equally sized subregion. A simple example goes back to Reichenbach (1956)—if the weather has been hot for the last few days, it is more likely that the weather will be hot tomorrow than that it will be cold, even if it is assumed that, in the long run, hot and cold temperatures are equally frequent. In this way, knowledge of the recent past can have a bearing on our choice of probability measures for the present state of a system.<sup>15</sup>

This is not to concede, however, that we are justified in using a uniform probability measure in cases in which we have absolutely no grounds for making any claims about the past of a system. For if we truly have no grounds for making any claims about the past of a system, then we are surely not justified in describing the present state of the system with *any* probability measure.

**6. Conclusion.** Although there are undoubtedly other arguments for

14. Another serious problem is that it is not clear why we should assign probability 0 to a set of measure 0. See Sklar 1993, 182–188, for a good discussion of this point.

15. This objection seems closely related to point (3) of Sklar 1993, 159.

Probability Principles, I conjecture that they all fail for similar reasons. There is no way to account for the variety of knowledge we might or might not have about the past of a system with a simple, a priori, principle that focuses solely on the present macrostate of the system. Once this is realized, some of the conceptual problems associated with statistical mechanics can be resolved.

## REFERENCES

- Albert, D. (2000), *Time and Chance*. Cambridge, MA: Harvard University Press.
- Callendar, C. (2004), "Measures, Explanation and the Past: Should Special Initial Conditions Be Explained?", *British Journal for the Philosophy of Science* 55: 195–217.
- Castell, P. (1998), "A Consistent Restriction of the Principle of Indifference", *British Journal for the Philosophy of Science* 49: 387–395.
- Earman, J., and M. Redei (1996), "Why Ergodic Theory Does Not Explain the Success of Equilibrium Statistical Mechanics", *British Journal for the Philosophy of Science* 47: 63–78.
- Reichenbach, H. (1956), *The Direction of Time*. Los Angeles: University of California Press.
- Sklar, L. (1993), *Physics and Chance*. Cambridge: Cambridge University Press.
- Van Fraassen, B. (1990), *Laws and Symmetry*. Oxford: Oxford University Press.
- Winsberg, E. (2004a), "Can Conditionalizing on the Past Hypothesis Militate against the Reversibility Objections?", *Philosophy of Science* 71: 489–504.
- (2004b), "Laws and Statistical Mechanics", *Philosophy of Science* 71: 707–718.