



ELSEVIER

Contents lists available at ScienceDirect

Studies in History and Philosophy of Modern Physics

journal homepage: www.elsevier.com/locate/shpsb



Laws and chances in statistical mechanics

Eric Winsberg

Department of Philosophy, University of South Florida, Fao 226, 4202 E Fowler Ave, Tampa, FL 33615, USA

ARTICLE INFO

Article history:

Received 27 July 2007

Received in revised form

3 March 2008

Keywords:

Thermodynamics

Statistical mechanics

Boltzmann

Lewis

Probabilities

Best system

ABSTRACT

Statistical mechanics involves probabilities. At the same time, most approaches to the foundations of statistical mechanics—programs whose goal is to understand the macroscopic laws of thermal physics from the point of view of microphysics—are classical; they begin with the assumption that the underlying dynamical laws that govern the microscopic furniture of the world are (or can without loss of generality be treated as) deterministic. This raises some potential puzzles about the proper interpretation of these probabilities.

© 2008 Published by Elsevier Ltd.

When citing this paper, please use the full journal title *Studies in History and Philosophy of Modern Physics*

1. Introduction

Statistical mechanics (SM) involves probabilities. At the same time, most approaches to the foundations of SM—programs whose goal is to understand the macroscopic laws of thermal physics from the point of view of microphysics—are classical; they begin with the assumption that the underlying dynamical laws that govern the microscopic furniture of the world are (or can without loss of generality be treated as) deterministic. This raises some potential puzzles about the proper interpretation of these probabilities. It also raises, more generally, the question of what kinds, if any, of objective probabilities can exist in a deterministic world.

E-mail address: winsberg@cas.usf.edu

The approach to the foundations of SM that is currently enjoying the greatest popularity, for example—the one that is sometimes referred to as “Boltzmannian,” since it represents a development of ideas originally put forward by the physicist Ludwig Boltzmann—treats the universe as a classical, deterministic system governed by Hamiltonian mechanics, and it postulates a probability distribution over a set of possible initial conditions for the universe. But there is only one universe, and it began only once. *Prima facie*, at least, this makes the fundamental nature of those probabilities somewhat mysterious.

It is a fairly straightforward corollary of most interpretations of probability that in a world governed only by dynamical laws that are deterministic, there can be epistemic probabilities—probabilities that reflect various sorts of degrees of belief—but no objective chances in the world. Recent arguments, on the other hand, with a growing group of adherents, purport to show that the probabilities in the Boltzmann program *cannot* be understood as *degrees of belief*. (Albert, 2000; Loewer, 2007)

If these arguments are sound, and such probabilities must be understood to be about objective chances—chances that are, as opposed to mere degrees of belief, part of the furniture of the world—then we are faced with the puzzling possibility that there can be objective chances in a deterministic world. What are we to make of this?

In what follows, I will explore these questions. In Section 2, I begin by reviewing the role that probabilities play in SM, and I survey the arguments in favor of the claim that relevant probabilities have to be objective chances. But most traditional accounts of objective chance have trouble making sense of non-trivial chances in a deterministic world. In Section 3, I review some of these accounts, including frequencies, propensities, and David Lewis’ “best system” account of chances, and highlight these difficulties. The fundamental problem with understanding the probabilities in SM to be objective is that we are meant to posit a probability distribution over a set of possible initial states, while we suppose, at the same time, that in fact only one of these initial states actually obtained.

In Section 4, I begin to tackle the problem by positing what I call a metaphysically minimal account of the probabilities. On the minimal account, the probabilities in SM are a convenient shorthand for summarizing non-probabilistic facts about the initial state of the universe. In Section 5, I argue that Barry Loewer’s recent proposal—a modified version of David Lewis’ Humean account of objective chance—is a metaphysical extension of this minimal account. Loewer’s account, I argue, is more “metaphysically thick” than the minimal account because Loewer is committed to the idea that the probabilities in SM are not “mere degrees of belief.” In fact, what Loewer is doing (and here he follows Lewis) is trying to construct something that mimics metaphysically inflationary non-Humean chances but which is reducible to the Humean mosaic. In Section 6, I examine, and subsequently reject, this Lewis/Loewer proposal. I argue that there is an insurmountable tension between the metaphysical thickness with which Loewer wants to endow his probabilities and the modifications he needs to make to Lewis’ original system in order to accommodate the needs of SM. So finally, in Section 6, I go back to the minimal account of Section 4 and give it a more deflationary and, I would argue, more strictly Humean interpretation. I argue, in the end, that the probabilities in question are, after all, degrees of belief.

2. A statistical hypothesis

Here is the Boltzmannian program in a nutshell. We begin with an empirical generalization. When two bodies that are not in equilibrium with respect to each other are brought into thermal contact, the combined body tends to approach thermodynamic equilibrium. This is a macroscopic observation. The Boltzmannian would like to explain this from the point of view of microphysics.

We would like a global account, so we begin by representing the complete state of the world, assumed to be a system made up of N classical particles, by a point \mathbf{X} in state space. The point \mathbf{X} is denoted by $(\mathbf{x}_1, \mathbf{p}_1, \mathbf{x}_2, \mathbf{p}_2, \dots, \mathbf{x}_n, \mathbf{p}_n)$, where \mathbf{x}_i and \mathbf{p}_i are the position and momentum, respectively, of the i th particle. We then define the notion of the “volume” of a macrocondition—an observable condition realizable by potentially very many different microscopic conditions. If an isolated system of particles is in some macrocondition A , then we define the $6N$ dimensional volume of the

subregions r of the entire region R compatible with macrocondition A to be

$$|\Gamma_r| = \int_{\Gamma_r} \prod_{i=1}^N d\bar{x}_i d\bar{p}_i$$

normalized over Γ_R .

Next, we make a cosmological posit: that in the distant past, the universe began with a particular global macrostate of very low entropy (where entropy is proportional to the log of the volume occupied by the macrostate). This hypothesis is sometimes called the “past hypothesis” (PH) and the initial state it posits, the “past state.”¹ We then hope to show that the overwhelming majority of that volume is made up of thermodynamically “normal” microstates. Those are microstates of the universe which, as they evolve according to the fundamental microscopic dynamical laws, will lead the universe, and all of its energetically isolated subsystems, to display the kind of behavior that the laws of classical thermodynamics predict that they should.

This is the technical machinery of the Boltzmann program. There are worries about this side of the problem. There are worries, for example, that it is plausible to suppose that any suitable candidate for the actual Hamiltonian of the universe is friendly enough to grant us what the Boltzmannian needs to conjecture to be the case. There are also worries about whether the Boltzmannian program is rich enough to deal with systems that are more complex than ideal gases. And finally, there are worries, given our current best theories of cosmology, about whether anything like PH is even a coherent supposition.²

But let us set those worries aside. Let us assume that something like the above is true. Suppose we assume it to be true that the overwhelming majority of the volume of state space compatible with the past state is made up entirely of microstates that lead, via whatever are the true microscopic laws of evolution, to macroscopic behavior that almost everywhere is compatible with all of the empirical generalizations that we have labeled “the laws of thermodynamics.” Suppose we knew that claim, let us call it “the Boltzmann Conjecture,” to be true, and suppose we accept something like PH. What then, in the form of a claim about probability, would we need to add to this technical machinery in order to explain and predict the past and future thermodynamic behavior that we are inclined to explain and predict?

One view is that the “Boltzmann Conjecture” is all that we need. According to Sheldon Goldstein, what the Boltzmann Conjecture tells us is that “*typical* phase space points yield the behavior that it was our (or Boltzmann’s) purpose to explain.” (Goldstein, 2001, p. 52) It follows from this, Goldstein claims, that “we should expect such behavior to prevail in our universe.” “What more,” he asks “could reasonably be expected by way of explanation?” (p. 52)

But surely we can, and should, expect more. We want to know what will happen in our world, not in a typical one. Indeed, Goldstein admits that as a physicist, he is satisfied with this answer, but “as philosophers we might in fact be inclined to demand more.” (p. 52) Indeed, as natural as the volume measure we have defined might be, without, for example, some hypothesis or argument to the effect that the *typical* state, on that measure, is a *likely* state, then it is not at all clear that the Boltzmann conjecture, all by its lonesome self, can do very much work for us at all. What good is it, after all, to know that the “typical” microstate compatible with that past state yields the behavior we want to explain and predict, if we have no way of arguing that the actual original microstate of the universe is, or is likely to have been, a typical one?

What would appear to be required to make the Boltzmann strategy work is just that: something (either an empirical hypothesis or an a priori argument) that allows us to make the connection between typicality—in the sense of volumetric predominance in the set of possible initial conditions on the measure of volume we defined—and probability. The standard way to do this is to somehow *identify* that standard *volume* measure (sometimes called the Lebesgue measure) as the correct *probability* measure over the set of possible initial states. But just how we should do this is not exactly clear.

¹ See Albert (2000) for the recent canonical formulation of the Boltzmann program.

² See Winsberg (2004a) for the first worry, Uffink (2002) for the second, and Earman (2006) for the third.

There appear to be two ways that one might think this identification might occur. One is to think that we would need an added empirical hypothesis. The empirical hypothesis would be that there is a uniform probability distribution, on the standard measure, over the set of microstates compatible with the PH. Let us call such a hypothesis PROB. Postulating such a hypothesis would be one definitive way to identify volume with probability. The other is to think, to the contrary, that the connection between volume and probability can be had, so to speak, for free. The latter is to think, in other words, that the universe is overwhelmingly *likely* to have begun in one of the thermodynamically-permissible-behavior-yielding states simply in virtue of what we know about the universe and the fact that those states are typical.

The first approach—according to which we need an added empirical hypothesis—has, in the recent literature, been closely connected to the view that the probabilities involved are *objective*—as opposed to *epistemic*—chances. This is the view that the probabilities in SM are *part of the mind-independent furniture of the world*, as opposed to being *features of our knowledge of the world*. The proponents of this pair of views (and their close connection) have chosen as their foil the view according to which the identification of the volume measure as a probability measure is not an empirical hypothesis. According to this construal, the identification is supposed to lack added empirical content. The entirety of the content of the probabilistic hypothesis is supposed to be motivated simply by reasoning about what we mean when we say that we know that a system is some macrostate, and that we are ignorant of the particular microstate through which that macrostate is instantiated. What makes the volume measure a probability measure, in other words, is simply the fact that it is a “natural” measure, and that we are uniformly ignorant of the particular microstate in which the universe began, and each equal volume region is therefore, in the epistemic sense of probability, equally likely to obtain. This view has come to be closely connected, in this same literature, to the view that the probabilities in SM are epistemic—that they reflect something about our degrees of belief—either that we have or that we should have—about how the world is.

Though both of these views—the view that the desired measure is without empirical content, and the view that probabilities in question are not about degrees of belief—have drawn fire from the same quarters (Albert, 2000) we should keep them, and the arguments against them, distinct.

Against the first view—the view that an added empirical hypothesis is not needed—David Albert argues that it is very difficult to see what role my uniform ignorance about the initial state of the universe could play in understanding the behavior that I want to explain. The universe, after all, begins in whatever microscopic initial state it begins in, will evolve in accordance with its laws. It cares not one whit about about my uniform ignorance. Albert writes: “Can anybody seriously think that it is somehow necessary, that it is somehow a priori, that the particles that make up the material world must arrange themselves in accord with what we know, with what we happen to have looked in to” (2001, p. 64).

I take this to be a rather forceful argument, so I think that we should focus on the alternative view. Whatever the identification of probability comes from, I take the Albert argument to be a rather compelling argument that it has empirical content—that it is not, so to speak, something innocent, something that we know a priori.

Against the second view, Albert argues as follows: “Can anybody seriously think that our merely being ignorant of the exact microconditions of thermodynamic systems plays some part *in bringing it about*, *in making it the case*, that (say) milk dissolves in coffee? How could that be?” (Albert, 2000, p. 64, my emphasis) In a similar vein, Barry Loewer argues: “PH and PROB *can bestow lawfulness* on [the higher level laws] only if they themselves are laws and if PROB is a law then the probabilities it posits must be objective; *not merely degrees of belief*.” (Loewer, 2007, pp. 304–305, my emphasis)

Let us accept, therefore, that the first argument shows that PROB must be an empirical hypothesis. And let us entertain, for the moment, the idea that the second set of arguments shows us that the probabilities in PROB are not degrees of belief, but instead summarize mind-independent facts about the world. If this is correct, what could the probabilities in such a hypothesis be about? There is, after all, so far as we know, only one universe, and it only began one time. And it began, we would like to suppose, not only in some particular macrostate, but in some single particular microstate. Indeed, it is plausible to think that, on the traditional picture of how classical physics works, there are only two sorts of basic facts about the universe: The fundamental (microscopic) laws that govern its evolution

on the one hand, and its initial (microscopic) state on the other. What sense does it make, then, to say that there is an objective probability distribution over a set of *possible* microconditions in which it began? We are forced to ask: what is the nature of such a hypothesis?

3. Traditional accounts of probability

We are looking for a plausible interpretation of the probabilities in PROB according to which they are objective chances. To find one, we need a conception of probability whose truth conditions are such that they can be satisfied by the facts that we are willing to posit about our universe in a mind-independent way. That is not as easy as it might sound.

Clearly, a simple frequentist account of probability will not work, as the origin of the universe is a singular event. There is only one universe, and it began only one time. But the problem is not that simple. There are, after all, two traditional accounts of probability available for meeting that need—there are two objective conceptions of probability that are capable of handling single instance probabilities. One is David Lewis' Humean account of "chances," and the other is the propensity account of probability.

Lewis' theory of chances, as he initially conceives of it in Lewis (1986), cannot be applied in a straightforward way. That's because Lewis developed his theory of chances as a central component of his defense of a thoroughgoing Humean supervenience. This means that, on Lewis' view, any two worlds that have the same spatio-temporal distribution of local properties have the same chances. Chances, in particular, arise when the "best" analysis of the regularities among those properties gives rise to chance laws: laws that postulate fundamental objective probabilities. For Lewis, the only chances come from the laws, which in turn come from the best analysis of the regularities found in the distribution of local fundamental properties.

So it would appear that in a world whose best system analysis (BSA) gives rise to a set of deterministic laws, there can be no objective chances as Lewis defined them. But the Boltzmannian approach is supposed to be able to explain how *deterministic* microlaws can give rise, statistically, to almost-everywhere-thermodynamic behavior. So if we assume, as I will do for the moment—much more on this later—that the law-like axioms of a Lewis best system are the microlaws of evolution, then there can be no Lewis chances (chances precisely as Lewis defined them) in classical Boltzmannian SM.

The traditional non-Humean rival to Lewis' account is the propensity interpretation of probability. It does not fare any better. The propensity interpretation identifies probabilities with frequencies of outcomes in hypothetical trials. As such, it can accommodate singular events, and it need not adhere to Humean Supervenience—two "Humeanly" identical worlds can have differing probabilities if they differ with respect to the frequency with which some type of event *would* occur in a hypothetical set of trials.

That sounds more promising. Still, the propensity account of probabilities only makes sense of probabilities of single instance events to the extent that they are the result of a (hypothetically) repeatable process or set of conditions.

On Karl Popper's account of the propensity interpretation of probability, for example, the probability of a particular outcome in singular instance event is equal to the relative frequency of that outcome in a virtual infinitely long sequence³ of trials "characterized by a set of generating conditions—by a set of conditions whose repeated realization produces the elements of the sequences. ... This means that we have to visualize the conditions as endowed with a tendency or disposition, or propensity, whose frequencies are equal to the probabilities." (Popper, 1959, p. 34).

In short, the propensity interpretation of probability can handle the probabilities of outcomes for singular instance events—and the probabilities need not even supervene on the Humean

³ It is actually not clear whether Popper's view of probabilities should be thought of as equal to frequencies in infinitely long virtual sequences, or approximately equal to frequencies in long but finite virtual sequences. Which of these two views one adopts has no bearing on this paper. For a thorough discussion of this point, and of the propensity interpretation of probability, (and in which the above quotation from Popper first appears) see Gillies (2000).

mosaic—but what is required even on the propensity interpretation is that the events arise as the result of some definite process under certain definite conditions—a process which, at least in principle, allows for any number of repetitions. In this respect, the propensity interpretation is no different than Lewis': it allows no non-trivial objective probabilities when there are only deterministic laws.

Thus, a proponent of this view would owe some account of how it makes sense to think of there being some process, under some set of conditions, which led to the initial conditions of the universe. But the very idea of this appears to contradict the definition of "initial conditions of the universe." The process and its conditions would have to lie, in some sense, outside the history of the universe. We would have to imagine something like a god who spins a cosmic roulette wheel in order to decide in which initial state to place the universe. It is not at all clear what sort of story could be told here, and the problem seems like it might be insurmountable.

A small digression can illustrate the severity of the problem. Consider the kind of worry advanced by Huw Price (1996) about the likelihood of the low-entropy past state. Recall that in order to get off the ground, the Boltzmannian approach requires a cosmological assumption: that the universe began in a very-low-entropy state. But if we are inclined to believe that the standard volume measure is a probability measure, then Price wonders how we are going to explain the fact that the universe began in what, on that very measure, is an incredibly unlikely initial *macrostate*—one whose volume is incredibly small.

A clear response to this worry has come from Craig Callender (2004). Callender argues that Price's worry is misguided, because the initial state of the universe is a boundary condition, and we need not be responsible for explaining boundary conditions. Callender gives the example, due to an exchange between Hugh Mellor and John Leslie, of a firing squad of 50 men aiming at you and all missing. Do we owe an explanation? Only, says Mellor (2002), if you believe in the ability and intentions of the shooters and in the accuracy of their weapons as causes of the bullets' trajectories. Not so, on the other hand, if you believe that the universe came into being right before the bullets go by you. In such a case, there is no mechanism to direct the bullets at you, and no need to explain the failure of any of the bullets to hit you. If the initial trajectories of the bullets have no mechanism behind them, argues Mellor, then they are just boundary conditions on the universe, and require no explanation.

Analogously, in his reply to Price, Callender argues that the low-entropy initial state of the universe is not in need of any explanation precisely because there is no mechanism that brings about the past state (simply in virtue of it being the past state). It is precisely the lack of such a mechanism that makes Callender's reply to Price viable (and I think correct). What this illustrates, though, is that anyone who wants to adopt a propensity interpretation of the probability distribution in SM will face a dilemma. To make literal sense of the probability distribution, (at least on the propensity interpretation) one needs to envision some sort of mechanism, or process, or set of conditions (with probabilistic laws to guide it), that brings about the initial state of the universe. And as difficult as that sounds, success would only plunge you directly into Price's trap—having to explain why that process led to such an unlikely outcome: a very low-entropy past state.

4. A minimal account

It seems clear at this point that something puzzling is going on. We are making the claim that PROB is an empirical hypothesis, but it is not clear what kind of hypothesis it could be. Such a puzzle produces the following temptation: let us strip PROB down to the most metaphysically minimal interpretation we can imagine. In that spirit, we might approach the problem along the following lines: let us forget, for a moment, about what we think PROB really means, and focus on what kind of information PROB provides; on what we can do with it; on what inferences it allows us to make.

Imagine a scenario in which a physicist is having a conversation with God. The physicist, quite naturally, asks God to tell her what the universe is like. God begins by telling the physicist what the dynamic laws of the universe are. But of course this is not enough. Because she wants to be able to make predictions, the physicist asks to know in what particular initial condition the universe began.

To which God replies: “The universe has 10^{50} particles, how much time do you have?” Since the answer is obviously “not enough,” God suggests the following alternative: “Pretend *as if* there were a uniform probability measure over the set of microconditions that are compatible with whatever macrocondition in which the universe began. From that fact and whatever macroscopic observations you make, you should have enough information to make any practical macroscopic predictions that you like. That should save us some time.”

In other words, when we postulate PROB, what we are really postulating is that the universe began in some genuinely *particular* microcondition—but one which, because we have limited computational power, observational data, and space in which to write our hypothesis down, we can characterize with only limited precision. The only *genuinely ineliminable* facts about the universe, on this account, the only facts which god would have to set down, the only ones He would have to share with us if time were unlimited, are its dynamical laws and its precise initial condition. Only, in other words, the dynamical laws and the precise initial conditions are part of the genuine supervenience base of the world. All other information is redundant. Any talk about measures is just a convenient way of giving the limited information that we have about the initial state. As Goldstein argued, we know that the “typical” initial state leads to the behavior we observe. All we really need is to make the empirical hypothesis that the initial state was in fact “typical.”

According to the minimal account, we talk about a probability measure only because talk of such a measure is necessary to allow us to make such a claim. And making such a claim allows us, with greater ease, to make some of the same predictions we would make if we had infinite computational power and access to the initial state of the universe in all of its gory and microscopic detail. Let’s call this the *minimal interpretation* of the probabilities in PROB.

Note that nothing that has been suggested falls under the scope of Albert’s argument. On the minimal interpretation, the identification of volume with probability is not taken to be without empirical content. Nor is it meant to describe a uniform ignorance about the initial state. It is meant to be a short hand for summarizing what in the way of empirical knowledge we have surmised about the initial state.

In fact, I take it that this is one way of reading David Albert’s own exposition of the Boltzmannian program when he states the content of the statistical hypothesis in these terms: “The *right probability distribution to use for making inferences* about the past and the future is the one that’s uniform, on the standard measure, over those regions of phase space which are compatible with whatever other information—either in the form of *laws* or in the form of *contingent empirical facts*—we happen to have.” (Albert, 2000, p. 96, first emphasis mine, the others Albert’s).

Let us now proceed in the following way: let us suppose that whatever detailed metaphysical story we tell about what kind of facts PROB purports to tell us, that this metaphysical story will honor the core of the minimal story—it will honor the idea that whatever PROB actually is about, the work that it does is the work that is described by the story about God and the physicist—it will honor the idea that the genuine non-redundant fact about the world is the actual initial microstate, but that PROB is the right hypothesis to use for epistemic agents like us (or for deities who want to communicate with us in limited time).

It is tempting to stop here. After all, the minimal account tells us that PROB is “the right probability distribution to use.” On the minimal story, God has given us the right (i.e. objectively correct—after all, He is God) distribution to guide our credences about future and past events, given our macroscopic knowledge. David Lewis’ “principal principle” teaches us that objective chances should turn out to be precisely those things that rationally guide credences, and on the minimal account, PROB seems to do this. So why not stop here?

It is time to clear up an ambiguity regarding the term “objective” in the expression “objective chance.” Objective probability is sometimes contrasted with the idea of subjective credences. In other words, the concept of objectivity involved is the concept whose opposite is subjectivity. But there is also a concept of objectivity a close synonym for which is “real”—as in “really part of the furniture of the world.” On that conception of objective chance, objective chance is contrasted with epistemic probabilities. It is often assumed that non-subjective probabilities are part of the furniture of the world, and that epistemic probabilities are necessarily subjective, but there is nothing logically necessary about this.

When Loewer argues that “mere degrees of belief” are insufficient to bestow lawfulness on the higher level laws, and when Albert says that what is required of PROB is something that can “bring things about” and “make things the case” it is clear that they are looking for “objective probabilities” in a sense which goes beyond the idea of being “objectively correct.” For them, clearly, the minimal account is a bit too minimal.

If Albert and Loewer’s arguments are correct here, then what we would seem to need is a story—a metaphysical story—about probabilities, which provides grounds for us to understand what sense, if any, there is to the idea that PROB is about objective chances. A story where “objective” is understood not merely as the opposite of “subjective,” but is understood to mean “really part of the furniture of the world,” rather than having anything whatsoever to do with any particular agent’s degrees of belief. To avoid confusion, let us, from now on call these kinds of probabilities “ontic” instead of “objective.”

Now, it might be objected that the idea of ontic probabilities is a bit obscure for a Humean like Loewer. What, after all, does it mean to say that a probability is “really in the world,” if after all, as a Humean, what you really believe is in the world is nothing more than the Humean mosaic. A couple of things can be said about this. First, what I mean to signal by calling a certain kind of probability “ontic” is that it meets *whatever* Albert and Loewer have in mind when they argue that the relevant probabilities are not “mere degrees of belief.” I mean to signal that they are the sorts of things that are plausible candidates for “making things the case” and “bringing things about” whatever those sorts of things turn out to be. Second, I presume, at least in the case of Loewer, that the idea of a probability being more than a “mere degree of belief” is associated with the Lewisian idea of a certain kind of Humean *reductionism*. I will have more to say about the idea of reductionism in a moment, when I return to Lewis. But for now, suffice it to say that it seems as though, for Lewis, something gets to count as reduced, if it recaptures—in Humean terms—the salient features of its metaphysically robust, non-Humean counterpart. (In much the way that Galilean free-fall is reduced to Newtonian physics if everything about the law of free-fall can be recaptured from Newton’s laws.) And it is to this standard that I will hold the Humean idea of a probability that is ontic, (that is, not merely a degree of belief): it has to recapture, by way of Humean reduction, the salient features of a metaphysically robust notion of probability. This will be an important point as we proceed.

5. A Humean answer

So what can be done about this? Can we, for example, deliver on the idea of understanding how our minimal account could be pointing at a story in which there are ontic probabilities in the world? We saw that, on the Lewis conception of objective chances, there was little hope of interpreting an initial probability distribution so long as we took the laws of the best system of the world to be the deterministic laws of evolution. But what if we take PROB itself to be a law? Why could we not suppose that according to the best system, the laws of evolution for the universe are deterministic, but that it is also a law that there is a uniform probability distribution over the set of microstates compatible with the past state?

This is the suggestion that has been made by Barry Loewer. It is worth quoting him at some length here:

What is the metaphysical status of PROB and PH? There is a great deal of controversy concerning how to understand probabilities in SM. The usual view is that they are ‘ignorance’ probabilities; i.e. the degrees of belief that one ought to have concerning the initial conditions of a system given that one knows that the system satisfies certain constraints (e.g. the system’s temperature). There is little discussion of the status of PH but it is generally construed as a *merely contingent* (i.e. not law-like) statement. I think that both of these interpretations are mistaken. PROB and PH (together with the dynamics) are best considered to be laws and the SM-probabilities to be objective. The reason is that they underwrite many of the asymmetric generalizations of the special sciences especially those in thermodynamics and these generalizations are considered to be laws. PH and PROB can bestow lawfulness on them only if they themselves are laws and if

PROB is a law then the probabilities it posits must be objective; not merely degrees of belief. (Loewer, 2004, pp. 304–305)

But what can it mean to say that the probability distribution is itself a law? In principal, such a suggestion might be open to a proponent of any sort of metaphysical picture of laws. Loewer himself, however, is quite clear that he favors a Lewis/Hume account of laws.

It is interesting that Lewis' Humean accounts of laws is friendly to the proposal that PH and PROB are laws. ... According to Lewis a law is a contingent generalization entailed by the Best System of the world. The Best System is that true theory couched in a vocabulary that includes a conditional probability function, mathematical notions, and terms that represent fundamental properties (Lewis' calls them 'perfectly natural properties'). What makes it *BEST* is that it best combines simplicity, informativeness and fit. Adding PH and PROB to the dynamical laws results in a system that is only a little less simple but is vastly more informative than is the system consisting only of the dynamical laws. (p. 305)

Loewer's proposal is that the probability distribution is a law precisely because it is part of a system that provides us with the most simple, fitting, and informative way of organizing what we know about the universe. As such, Loewer's proposal is precisely an attempt to do what I described above—to recapture the idea of a metaphysically robust probability from the Humean mosaic. It is an attempt to metaphysically soup up the minimal account.

For a Humean, the only fundamental facts are the facts about that mosaic. To a Humean then, it is clear that PROB would necessarily be redundant—since to a Humean, a full specification of the Humean mosaic entails a full specification of all the genuine facts about the world. And so any Humean will agree with the minimal account. But the souping up that Loewer is trying to do is to effect the kind of reduction I spoke of above—it would turn our minimal account of PROB into an account in which the probabilities in PROB (in virtue of PROB being a law) become genuinely mind-independent probabilities. What Loewer seems to be hoping to get is something that preserves all of the elements of a metaphysically thick notion of probability (for example: the propensity notion), but be able to completely reduce it to the Humean mosaic.

The trick has the hope of working because in Lewis's framework, there's nothing more to being genuine, in the modal realm, than occupying the appropriate place in the best theory—it is a very metaphysically deflated conception of laws and chances, compared to typical non-Humean views. The Humean can insist that the only fundamental thing in the supervenience base of the world is its initial microstate, but also insist that the probabilities in PROB are as much “part of the world” as rocks, snails, and sporting events. This is in sharp contrast to non-Humean, who, as we have seen, has a very hard time explaining how there could possibly be some fact of the matter about the initial state of the universe, and some mind-independent fact about non-trivial probabilities for such various possible initial states.

Still, we need to keep distinct Loewer's view from the minimal account of which it is an extension. On the first view, we have simply highlighted the fact that, from the point of view of making certain macroscopic predictions, PROB is a useful shorthand substitute for the actual initial microcondition of the universe. On Loewer's view, PROB is still shorthand, but it is not just *any old shorthand*, it is the shorthand that is part of the best system, and hence is a law, and hence is supposed to underwrite a Humean reduction of genuine ontic chances. Hence the Loewer proposal is one that starts with something like the minimal view, but tries to provide a unified metaphysical story about laws and probabilities according to which the probabilities are able to take on the role that Albert argues that they need to take on: to *explain* why the ice melts in my cocktail, and to *make it the case* that the milk mixes in my coffee.

Loewer's proposal is, as such, very much in the original Lewisian mold of Humeanism: the project of preserving the fundamental features of a non-Humean metaphysics—arguably: to preserve things that can do what “mere degrees of belief” cannot do—while at the same time reducing them to a supervenience base of the Humean mosaic.

But to succeed in this project with regard to the probabilities in SM, the proponent of the Lewis/Loewer scheme has to make three substantial changes to Lewis' program before she can pronounce

PROB to be a genuine chance law and, in so doing, give a principled metaphysical account of how these probabilities could be ontic—of how they could explain why the ice melts in my cocktail and underwrite the laws of the special science. In what follows, I will argue that these changes destroy the plausibility of the proposal doing the work that Loewer would like it to do—of recapturing ontic chances in a framework that reduces to the Humean mosaic.

As we go through these arguments, we need to keep our eye on that very important ball: the distinction between ontic chance, on the one hand, and objective chance, in the sense of “objectively correct epistemic probability,” on the other. We have to remember the point of this whole exercise. The original arguments, put forward by Albert and Loewer, were meant to show that the probabilities of SM cannot be epistemic probabilities, because there is no way that my state of mind can explain why the ice melts in my cocktail—and that epistemic probabilities (“mere degrees of belief”) cannot do the work of underwriting the laws of the special sciences. And so the relevant test is to see whether Loewer’s account can successfully deliver on *this* promise: to make sense of the idea that PROB, in addition to being “the right probability distribution to use” in fact mimics the behavior of more ontologically thick probabilities.

Indeed, if Carl Hoefer is correct, Lewis was “trying to offer an account of objective chance that mimics the way we think of chance when we think of them as *propensities*, making things happen (or *unfold*) in certain ways.” (2007, p. 6) I think that Hoefer is indeed correct in his diagnosis of Lewis’ intentions, and I think this highlights an important respect in which Lewis was trying to create a competitor for propensities—that is to say an account of ontic chance. It is in this sense that I think it is correct to think of Lewis as wanting to *preserve* all of the features of propensity chances and then *reduce* them to the Humean mosaic.

Lewis, for his own part, took a “spoils to the victor” line about his BSA of laws and chances, and explanation. He believed that his BSA delivered the genuine laws and chances, and that these laws and chances, in turn, give you all of the explanatory power you want for free. Let us accept that this was an acceptable line for Lewis to take. And to the extent that we accept this line, let us say that, for all intents and purposes, Lewis offered us an ontic account of chance.

But we need to be very careful. Because it may turn out that the details of the account Lewis offered matter to whether or not this “spoils to the victor” line remains plausible. We will have to look at the details as they come, but for now, we can certainly note one feature of Lewis’s scheme that seems crucial: Lewis proposed a story on which the BSA results in one unique answer—what we come up with is *the* best system of the world—and hence there is always one unique chance of any particular event. This is a crucial difference between an account of objective epistemic chance, and one that purports to be an account of ontic chance. It is plausible to suppose that what is an objectively correct epistemic probability of a single event for an agent in one epistemic situation might very well differ from the objectively correct probability of that same event for another agent in another epistemic situation. It is plausible to suppose that an account of objective epistemic chance could be based on a BSA in which what counts as “best” varies according to the epistemic agent doing the analysis. But a single event can only have one ontic chance. Only *the* chances can do explanatory work. If spoils go to the victor, they have to go to *the* victor—they cannot go to the victor under such-and-such circumstances.

In other words, it seems to me that whether or not Lewis intended his account to be about ontic chance, or about objective epistemic probability, there are certain features of his account that seem designed to make it plausible, and for which I would argue they are crucial, that the account delivers chances that can play the role that ontic chances are ordinarily supposed to play—that they be able to do what Albert and Loewer think degrees of belief cannot do. And so these are also then features, *inter alia*, that are crucial to making it plausible that one could hold to the “spoils to the victor” line at the end of the story. The import of this will become clear shortly.

In what follows, I argue that there are three such features of Lewis’ account, and I argue that these are features that are abandoned by Loewer.

First, in Lewis’s BSA: “The laws are those *regularities* that are theorems of the best system.” (Lewis, 1994, p. 480) In his primer on Lewis, Loewer himself wrote: “The laws are the *generalizations* entailed by the best system for [a given] world.” (1997, p. 106)

But surely not every axiom, nor every theorem, of such a system would necessarily be a regularity or a generalization. Even if PROB and PH are part of the best system, it is not obvious that they should count as laws according to Lewis' scheme, since neither one is a regularity or generalization. PH is clearly a singular fact. And to the extent that PROB stands in—in simpler form—for a description of the initial microstate, it is also certainly not a regularity, and it is hard to see how it counts as a generalization.

Of course we could monkey with the requirement that laws be regularities. We could even rephrase PROB so as to make it sound like a regularity: "For any time t , if the universe's macrostate is Mt , then the probability that its microstate is in A (subset of Mt), is given by the volume of that region of A that is dynamically compatible with the past hypothesis." (Call this PROB-R) But that does not seem to be the kind of regularity that Lewis was after.

Why not? If I am right that we should understand Lewis as a reductionist about chances—as someone, that is, who was interested in preserving all of the features of propensity chances in his account, than he would have been interested in regularities that unfold in time. This is, after all, a feature of propensity chances. Correlarily, to the extent that Albert and Loewer claim to be after probabilities that can do things that epistemic probabilities cannot do—make things the case, bring things about, and underwrite laws—they too would appear to be committed to offering a reduction of a metaphysically thick conception of probability to the Humean mosaic. Therefore, they too, perhaps, ought to be committed to regularities that unfold in time.

Of course, one could perfectly well argue that the metaphysically thick concept of chance to be reduced need not be (as Lewis seems to have supposed) exactly the propensity concept. There could, presumably, be a suitably thick conception of chance that could do the work and not be the propensity conception, and therefore, not necessarily unfold in time.

I think, therefore, that we can set this first worry aside. Just because Lewis felt the need to have his chances exactly mimic propensities does not mean that any genuine candidate for the proper conception of chance to be reduced needs to do the same. But, I would argue, there *are* some essential features of propensity chances that need to be preserved in the reduction if one is to succeed in reducing something about which one could plausibly say it could *do* certain things that epistemic probabilities cannot do.

This brings us to our second point. The kind of informativeness Lewis expects chance laws to have is that they should tell us "what the chances are of various outcomes in various situations" (p. 480). In order to play this role, chances of events are derived from the "history to chance conditionals" (HCCs), which are in turn entailed by the laws. Again, it is useful to quote from Loewer's primer:

The totality of chance laws entail what Lewis calls the "history to chance conditionals"; statements of the form $H_t \rightarrow \text{Pr}(E) = x$. These specify the chances of future courses of events after t if the history up through t is H_t . The chance of E at t is derived from the history up to t and the history to chance conditionals. (Loewer, 1997, p. 106)

Unfortunately, there is some ambiguity here. Loewer claims that PROB (or PROB-R) is supposed to add to the informativeness of the microlaws by giving the chances of macroscopic events. But how, in Loewer's story, are we supposed to understand PROB-R giving us a chance of event E after time t that is non-trivial—that is not zero or one? Presumably, the story goes something like this: We encounter some system (say, a cocktail) at time t and we want to know the probability of some event E (say, the ice melting) occurring. H_t , the history of the world up to time t , will entail what macrostate the universe is in at t , and then that macrostate at t gets plugged into PROB-R to entail probabilities for future macro-events E at different times using the microlaws. So PROB-R entails statements of the form: $H_t \rightarrow \text{Pr}(E) = x$, where x can be anything from zero to one. This looks good.

But hold on. The ambiguity I mentioned is about what kind of information is meant to be contained in H_t , because if H_t contains all of the microscopic information, that is if H_t includes a specification of the microstate of the world at time t , then H_t plus the microlaws will either entail E or not E , and so the probability of E will always be zero or one. And without the microlaws, chances of future events are undefined. If all of the HCCs in Loewer's account give chances of events that are zero or one, then we are in trouble.

So we should ask whether the information contained in H_t should be restricted to macroinformation. Lewis would say “no.” He would have us include in the history of the world up to t all of the microscopic information about the world. So if Loewer wants PROB to generate non-trivial chances, he will have to alter this feature of Lewis’s account. In fact, he is going to have to alter Lewis’ account of “admissibility.” He will have to claim that information about the microstate is “inadmissible” because otherwise the principal principle will be violated—as any rational agent who had specific information about the microstate would rationally assign himself a credence that differed from the chance that Loewer would have us calculate (the one you get from limiting H_t to macroscopic information).

One might of course ask why Lewis was so interested in chances defined in terms of HCC, rather than future-to-chance, or left-half-of-universe-to chance, or whatever. Similarly, we might also ask why the information that gets plugged into the notion of “the history of the world up to t ,” that is part of the idea of a HCC, has to be all of the *microscopic* information, rather than, say, the *macroscopic* information. Why should all microscopic information be admissible?

Presumably this is because history-to-chance, and history-in-terms-of-the-micro, happens to be the one which relates most naturally to our idealized epistemic situation. This is the fundamental idea of the principal principle, and of the idea of “admissibility” that goes with it: objective chances are precisely those probabilities that would be equal to the credences we would assign if we were ideal epistemic agents with access to all the information that being an ideal epistemic agent entails. This accords perfectly with Hoefer’s diagnosis of Lewis as trying to mimic as much as possible of a metaphysically thick conception of chance like propensities.

Now, it is of course worth noting that this desire on Lewis’ part was perhaps in tension, if not outright contradiction, with his avowed Humeanism. Lewis’ Humeanism starts from the whole mosaic of history and builds up everything from that. In that mosaic, there is no “flow of time,” no “evolution” or “production of the future out of the past, by the governance of the laws.” All those metaphors should probably be rejected by a Humean.

Indeed, on Hoefer’s own Humean account of “objective chance,” those metaphors are rejected, and so is Lewis’ account of admissibility. So Hoefer can reject both this argument and the last one—for just this reason. Hoefer can reject the demand that the regularities that give rise to chances be the sorts of regularities that unfold in time. And he can certainly reject the demand that the relevant HCC be in terms of complete micro-histories. Most importantly, Hoefer can have an account of objective chance, and indeed he does, in which the chance of an event depends on how you construe the epistemic agent who is rationally guided by it.

But Hoefer’s is not an account of chance in which chances can push and pull in the world. His account is more rigorously Humean one. Hoefer, in short, is no reductionist about chances. Albert and Loewer, on the other hand, most certainly are. They want the kinds of chances that can do what, in Albert’s opinion, only a crazy person thinks that epistemic chances can do—*explain* why the ice melts in my cocktail, *make it the case* that milk mixes in my coffee, and *underwrite* the laws of the special sciences. So, avowed Humeans or not, like Lewis, Albert and Loewer need to offer, as the candidate for reduction, a conception of probabilities in which probabilities behave in ways that mimic things that can do these things.

And they need a conception of probabilities in which there are something more than “the right probabilities [for particular agents (for example, ones like us)] to use.” They need a conception in which probabilities are “the right ones to use” for one standardized sort of epistemic agent. That is what the principal principle needs to demand if it is meant to ground a metaphysically thick conception of ontic chance. And an idealized epistemic agent is surely one that reasons from complete micro-histories. And an idealized epistemic agent derives his HCC from the entirety of the laws available.

The moral of the story can be put rather concisely: the epistemic limitations of the agent under consideration can be part of the supervenience base in a notion of objective epistemic chance, but they cannot be part of the supervenience base of ontic chances—of things that cause ice to melt in my cocktail. (To paraphrase Albert: “Can anyone seriously think that my own epistemic limitations play some part in bringing it about, in making it the case, that (say) milk dissolves in coffee? How could that be?”)

There is another important point to make about admissibility. On Lewis' account, all of the microfacts about the world up to time t are admissible. On Hoefer's account of objective chance, the criterion of admissibility is guided by the kind of chance rules that describe the chance setup in question. Any information that does not go into applying the chance rules in question is not admissible. To use one of his examples, if it is a chance rule that woman over the age of 40 have a $c\%$ chance of getting breast cancer, then for a chance setup under that description, it is inadmissible to consider information about a woman's genetic history. But then even on Hoefer's less metaphysically thick account, microfacts would be admissible in the case of chances being calculated using PROB. This is because, as we pointed out earlier, we get no predictions from PROB without employing the microlaws. And the microlaws operate, by definition, at the microscopic level of description.

This brings us to the third point: Lewis defined informativeness in terms of the fundamental language, what he calls "the perfectly natural properties." Loewer's proposal is to change that definition to include informativeness "relative to the thermodynamic language as well as the fundamental language" (Loewer, 2007, fn 23). But then it is not clear exactly what the game is supposed to be. It sounds as if it is supposed to be played in two steps:

Step 1. Define informativeness in terms of the fundamental language. Find the best system using that definition, and identify the regularities as the laws. These will be the microlaws.

Step 2. Keep the microlaws you found in step 1, but now redefine informativeness to include facts captured in the thermodynamic language. See what needs to be added to the laws you found in step 1. Call whatever you add a new law.

But this cannot possibly be a principled way of uncovering what the genuine laws of nature are, and hence it cannot be a principled way of uncovering what are the genuine objective chances. The problem is simple: the microlaws will be part of the best system if and only if informativeness is defined in terms of the fundamental language. That is because if informativeness is defined in terms of the thermodynamic language, then the best system will simply be the laws of thermodynamics. So you only get microlaws if you try to find the system that is maximally informative in the fundamental language. But it is only if we take it for granted that the microlaws will be part of the best system that PROB becomes a plausible candidate for being part of the best system when we try to maximize informativeness in the thermodynamic language. PROB does not fall out of any one-pass attempt to find the maximally best system. There is no single description of the world for which PROB+the microlaws is the best summary.

So PROB is very useful for making predictions in certain particular epistemic situations. But it is not part of any one single best system. So, again, it is implausible that PROB contains probabilities that are uniquely defined, objective, and principled enough to do what Lewis hoped his "chances" would do: mimic the status of something like metaphysically thick non-Humean chances such as propensities. But only if you do this mimicry is it plausible to think that they have effected a reduction of non-Humean concept to the Humean mosaic, and only then is the "spoils to the victor" line plausible.

The overall lesson is this: there appears to be a strong tension in Loewer's attempts to elevate PROB to a metaphysically thick law that can do explanatory work. Given Lewis' notions of "strength" and "simplicity," it is not clear that PROB adds the former at all; but if Loewer takes a stand on simplicity and strength that clearly lets in PROB, he seriously weakens his claim that the BSA laws *genuinely* explain, and he destroys the plausibility of the whole "spoils to the victor" argument. There is indeed a "best system" that includes PROB among its axioms, but it is a very pragmatic, user-friendly, elegant-summary kind of "best system," not a substitute for metaphysically "thick" and "fundamental" laws—what Lewis arguably was hoping for—and the kinds of thing that can give us ontic chances: things that can "make it the case," or "bring it about" (Albert) that ice melts in my cocktail. When Loewer's metaphysical thickness is the standard, the probabilities in PROB come up short.

We can conclude the following: a Humean will find something natural about our minimal account. But, despite the apparent connection between the minimal position and the view that PROB is a Lewis law, there is no good principled way (at least not one that follows Lewis even in general spirit) for her to make PROB out to be a law that underwrites genuine non-trivial ontic probabilities. And hence there does not appear to be a good, principled way to extend the minimal account into a robust Lewisian reduction of ontic chance in Boltzmannian SM.

6. Conclusions

What options do we have left for understanding probabilities in the foundations of SM? Below are some ideas.

To begin, we could relax some of the constraints on the requirements of counting a program as Boltzmannian. One way to do this is to look for the probabilities that we need in the dynamics themselves. If objective chances cannot be found in world with deterministic laws of evolution, then why not look elsewhere? The Boltzmannian program is generally assumed to be founded on a set of deterministic microscopic laws of evolution. We usually assume in the foundations of SM that if we could understand all the issues that arise in a classical, deterministic world, we could extend what we learn to our own, possibly indeterministic quantum world. And so most discussion of the foundations tends to ignore the possibility that the dynamics of our world might not be deterministic.

But on some interpretations of quantum mechanics (QM), the dynamical laws are inherently stochastic. Albert (2000) has made the suggestion that if there is the right kind of fundamental probabilities in our dynamics, we might be able to exploit those probabilities to solve some of the problems we have been talking about. While this option is certainly promising, it is also not without its costs.

To understand the cost, we need to remind ourselves that the reason we got into all of our trouble in the first place is that the Boltzmann conjecture (which was up until now we have been tacitly assuming was a conjecture about a set of deterministic laws) only promises that the overwhelming majority, on a particular measure, of the possible initial states of the universe will yield the behavior we want to explain. What we have been having trouble with is justifying or explaining our willingness to ignore that small minority of states—the “atypical” ones. To alleviate that problem by moving to a stochastic dynamics, what we would need to do is to replace the Boltzmann conjecture with a much stronger conjecture about that new dynamics. And it would have to be stronger in the following way: the “Super-Boltzmannian” (SB) conjecture, if you will, would have to be that every single last possible initial microstate of the universe that is compatible with the past state is overwhelmingly likely to lead, via those indeterministic laws of evolution, to everywhere thermodynamically permissible behavior. If there are any microstates that do not have this property, then we are no better off than we were in the deterministic case. But if the SB conjecture were true, then we would be free of our initial problem: trying to understand why the universe behaves as if it began in a “typical” state—since according to SB, every possible initial state is “typical.” The beauty of this approach, therefore, is that it would eliminate the need for an added empirical hypothesis in the form of PROB.

So the first cost of this approach is that we have to accept, without very much in the way of argument having to date been offered, the claim that the actual dynamical laws of our universe meet this rather stringent demand. On the other hand, of course, all our attempts to deal with the problems of SM are plagued by conjecture. On its face, the demand that the new SB conjecture apply to *every single possible initial state* does appear to make its plausibility a higher bar to reach. But sometimes it is more plausible to believe something about *every* member of a set, than it is to prove something about *one*, unknown member (namely, the actual world’s IC).

There is also a second, perhaps more serious cost associated with this proposal—the kind of cost associated with committing to any particular interpretation of QM. As Albert has argued, because the dynamics we would consider would have to ground a SB conjecture, only an “interpretation” of QM that has genuinely stochastic collapses of the wave function—one like the GRW interpretation—can do this kind of work. Bohmian mechanics, for example, clearly will not work.⁴ And like every interpretation of QM on the table, GRW faces some difficult obstacles. As long as those problems remain unsolved, hitching the solution to one foundation problem to another unsolved and deeply intractable one is obviously risky. If GRW happens to win out in the QM interpretation game, (and if

⁴ This is because the apparent probabilities in Bohmian mechanics are really the result of hidden states, and so if we wanted to rule out the “atypical” trajectories, we would have to rule out, or deem unlikely, certain initial wave functions for the universe. And this puts us right back where we started.

the SB conjecture turns out to be true), then this will certainly be a tremendous boon for the student of SM. But in the meantime, she is perhaps best off not betting everything on a flavor of QM that may not make it to the finish line.

Another constraint of the Boltzmann program that we could relax is the requirement that the story we tell be a global one. In other words, we have been assuming, up until now, that the reduction of the macroscopic laws of thermal physics to the microlaws take place on the scale of the entire universe, as opposed to the scale of each individually quasi-isolated system. If we limited ourselves to reducing the macroscopic behavior of such systems to the microlaws on a case-by-case basis, the task becomes comparatively simpler. If, for example, we suppose that something like PROB applies to each individual quasi-isolated subsystem at the moment of its isolation, then PROB takes on something of the character of a dynamical postulate, and most of the traditional interpretations of objective chance become viable candidates.⁵ But the costs of that sort of approach are quite steep. The principal cost, in particular, is that we have to give up on the idea that the microlaws offer us a complete coherent description of the world. The cost, in short, is fundamentalism.

So let us suppose, at least for the rest of this discussion, that we would like to stick within the broad constraints of the Boltzmannian approach. If, in that context, PROB cannot be about ontic chances, what sense would we make of it? At this stage, I would argue that that the right way to think about the problem is to return to the minimal account. The minimal account retains the core of Loewer's proposal. It retains the idea that PROB is, to quote Albert again: "The right probability distribution to use for making inferences about the past and the future. ..." It retains, furthermore, the idea that the only genuine facts about the world—in the sense of being part of the supervenience base of the world—that is built into PROB is the actual precise initial condition. It shares, with Loewer's proposal, the idea that PROB is a shorthand for our inferred knowledge about the initial state.

But what should we say about these "minimal" probabilities? Why do we seem to need to make use of probabilities in framing an *empirical hypothesis* if there are no genuine ontic probabilities in the Boltzmannian world? Why is the "right hypothesis to use" one that mentions probabilities if there are no genuine probabilities in the world? And what, in the end, "makes it the case" that the ice in my cocktail melts?

The answer, of course, has its roots in the limits, both computational and observational, as non-ideal epistemic agents. Recall that the Boltzmannian program rests on conjectures about the microlaws. For the same reason that we have little hope of proving these conjectures to be true, we also have no hope of reasoning directly from thermodynamic observations to specific details of the initial microstate. We can only reason statistically—that is, we can only use the kind of reasoning that we use to make the Boltzmann conjectures plausible.

We know that we need an empirical hypothesis about the initial microstate. The right empirical hypothesis to make is the one that encodes what we have surmised about the initial microstate of the world (from thermodynamic observations, induction about the future and past based on those observations, and Boltzmannian reasoning); one such that, when we combine it with what we know about a system of interest (its current macrostate) and apply the reverse of that reasoning, we get a certain (overwhelmingly high) probability of thermodynamic behavior. And so the best way to encode that particular hypothesis—the hypothesis that summarizes the great deal that we have surmised about the particular initial microstate that the universe, as a matter of fact, happens to have begun in—is with PROB. But PROB expresses neither an ontic probability about the initial state (we have seen that this view is untenable) nor a uniform absence of knowledge about the initial state (that will not do the work—and the Boltzmannian has in fact surmised a great deal about the initial state).

Perhaps this says enough. For some, this minimal account of the nature of PROB will suffice. PROB is the right probability distribution "to use" because it encodes, in a suitably useful way, what we

⁵ For an account of what such an approach to the foundations of SM would look like, and an explanation of why it requires us to abandon fundamentalism about the microlaws, see Winsberg (2004b). It is also worth noting that such an approach is cheaper from the point of view of what it requires us to conjecture about the microlaws. It has the potential to alleviate at least two of the three worries we mentioned, but dismissed, at the beginning of the paper.

have surmised about the initial state. Others, on the other hand, may want to know more. In particular, one might be inclined to ask the following question: if PROB consists of a probability distribution that “encodes our limited knowledge” of the initial state, why should we not say that the probabilities in PROB are epistemic probabilities (“mere degrees of belief”)? Furthermore, if PROB is “the right probability distribution to use,” then why not say that PROB is an objective(ly correct) epistemic probability distribution?

But recall the fact that part of what makes PROB the right probability to use is the set of our epistemic limitations. And notice the fact that there are plenty of other probability distributions we could use—ones, for example, that smear probability in lumpy ways around the past state—which would also make predictions that accord equally well with our experience. Part of what makes PROB the right probability to use is its mathematical simplicity. Hence the fact that PROB is the right one is a fact that supervenes on all (macro-) history, future and past—but that history includes, in part, *our epistemic limitations*, and the fact that PROB is the simplest and most mathematically tractable. So PROB is arguably part of the best system, but it is the best-system-for-us, not the best-system-for-an-ideal-epistemic-agent, and hence certainly not THE-best-system-of-the-world. So PROB can contain objectively correct probabilities (for us), but not ontic probabilities. They are objectively correct, epistemic probabilities.

But if all of this is right, and the probabilities in PROB are epistemic, what about the arguments by Albert and Loewer that epistemic probabilities are out of the question? There were two arguments to that effect:

Can anybody seriously think that our merely being ignorant of the exact microconditions of thermodynamic systems plays some part in bringing it about, in making it the case, that (say) milk dissolves in coffee? How could that be? (Albert)

PH and PROB can bestow lawfulness on [the higher level laws] only if they themselves are laws and if PROB is a law then the probabilities it posits must be objective; not merely degrees of belief. (Loewer)

But these arguments seem to rely on two assumptions. The first assumption is that if the probabilities in SM are epistemic, then the identification of volume with probability occurs a priori—then it is necessarily the case that we know it in virtue of the fact that we would come, *tabula rasa*, to the conclusion that we are uniformly ignorant about what microstate a system is in when we know it to be in some macrostate.

Some surrounding passages in Albert's (2000) make it clear this is how he understands the idea of SM being about epistemic probabilities. In characterizing the epistemic view, he writes that we identify volume with probability, on that account, simply in virtue of the fact that “nothing in our epistemic situation favors any particular one of the microconditions compatible with *X* over any other particular one of them.” And then he asks: “Can anybody seriously think that it is somehow necessary, that it is somehow a priori, that the particles that make up the material world must arrange themselves in accord with what we know, with what we happen to have looked in to.”

But while it is true that nothing in our epistemic situation favors any particular microcondition compatible with the past state, it is false that this epistemic situation is arrived at a priori. The particles have arranged themselves as they wish, and after much investigation of that arrangement, using macroscopic observation and statistical reasoning, we reach our epistemic situation.

The second assumption behind the arguments of Albert and Loewer is that there is a symmetry between “making things the case” and explanation, on the one hand, and prediction and retrodiction, on the other. Albert asks how our state of knowledge can bring it about, or explain, why cream mixes in coffee. It does not. What brings about the mixing, what explains it, is the actual initial state of the universe. What, about that initial state, that we actually know, is summarized in PROB. And so what I use to predict that the cream will mix is PROB. But that is not what makes it the case, nor what explains it.

Loewer's argument relies on a similar assumption. He wants to know how PROB can underwrite the laws of thermodynamics unless it is itself a law, with objective chances in it. But PROB does not underwrite any actual laws. Whatever actual macroscopic laws there are, they are underwritten by

the actual initial state of the universe. End of story. What underwrites my belief in the laws, my willingness to predict that these laws will continue to hold, my understanding of the laws of thermodynamics, is my epistemic state regarding the actual initial state of the universe. And that is summarized in PROB.

It is true, of course, that from the dynamical laws and the actual initial microstate, every true generalization follows—both lawful and accidental. And so microphysics will not be able to ground the difference, in the special sciences, between genuine laws and accidental generalizations. This is presumably part of what motivates Loewer. But why should we think it should? In the right spirit of Mill, Ramsey, and Lewis, the difference will be determined by deciding which generalizations are a part of the best system at whatever level of description they occur at. The second law is a genuine law of thermal physics, because, when the world is described in the thermodynamic language, the second law is part of the best system for summarizing those facts.

In sum, there can be objectively-correct-for-us epistemic probabilities in a Boltzmannian world, but there can be none of what we normally think of as objective (that is to say ontic) chances. And what probabilities there are do not “underwrite” the higher level laws, nor do they “make things the case” or “brings thing about.” Only microlaws and initial states can do that kind of work.

Acknowledgments

Thanks to Carl Hofer, Roman Frigg, Daniel Weiskopf, Steven Geisz, Mathias Frisch, Torsten Wilholt, and to my anonymous referees for helpful comments and advice.

References

- Albert, D. (2000). *Time and chance*. Cambridge: Harvard University Press.
- Callender, C. (2004). Measures, explanations and the past: Should “special” initial conditions be explained? *British Journal for the Philosophy of Science*, 55, 195–217.
- Earman, J. (2006). The ‘Past Hypothesis’: Not even false. *Studies in History and Philosophy of Science*, 37, 399–430.
- Gillies, G. (2000). Varieties of propensity. *British Journal for the Philosophy of Science*, 51, 807–835.
- Goldstein, S., (2001). Boltzmann’s approach to SM. In J. Brimont, D. Dürr, M. C. Galavotti, G. Ghirardi, F. Petruccione, & N. Zanghi (Eds.), *Chance in physics: Foundations and perspectives*. Lecture Notes in Physics 574.
- Hofer, C., (2007). The third way on objective probability: A skeptic’s guide to objective chance. *Mind*, 116(2).
- Lewis, D. (1986). A subjectivist’s guide to objective chance, *Philosophical Papers*, Vol. II. Oxford: Oxford University Press.
- Lewis, D. (1994). Humean supervenience debugged. *Mind*, 103, 473–489.
- Loewer, B. (1997). Humean supervenience. *Philosophical Topics*, 101–127.
- Loewer, B. (2004). David Lewis’ Humean theory of objective chance. *Philosophy of Science*, 71, 1115–1125.
- Loewer, B. (2007). Counterfactuals and the second law. In H. Price, & R. Corry (Eds.), *Causation, physics, and the constitution of reality: Russell’s republic revisited*. Oxford: Oxford University Press.
- Mellor, H. (2002). Too many universes. In N. A. Manson (Ed.), *God and design: The teleological argument and modern science* (pp. 221–228). London: Routledge.
- Popper, K. (1959). The propensity interpretation of probability. *British Journal for the Philosophy of Science*, 10, 25–42.
- Price, H. (1996). *Time’s arrow and Archimedes’ point: New directions for the physics of time*. New York: Oxford University Press.
- Uffink, J. (2002). Review essay: Time and chance. *Studies in History and Philosophy of Modern Physics*, 33, 555–563.
- Winsberg, E. (2004a). Can conditioning on the ‘past hypothesis’ militate against the reversibility objections? *Philosophy of Science*, 71, 489–504.
- Winsberg, E. (2004b). Laws and SM. *Philosophy of Science*, 71, 701–718.