

Chance in Boltzmannian Statistical Mechanics

Roman Frigg*

Philosophy of Science 75, 2008, 670-681.

Abstract

In two recent papers Barry Loewer (2001, 2004) has suggested to interpret probabilities in statistical mechanics as Humean chances in David Lewis' (1994) sense. I first give a precise formulation of this proposal, then raise two fundamental objections, and finally conclude that these can be overcome only at the price of interpreting these probabilities epistemically.

1 Introduction

Consider a gas that is adiabatically isolated from its environment and confined to the left half of a container. Then remove the wall separating the two parts. The gas will immediately start spreading and soon be evenly distributed over the entire available space. The gas has approached equilibrium. Thermodynamics (TD) characterises this process in terms of an increase of thermodynamic entropy, which attains its maximum value at equilibrium. The Second Law of thermodynamics captures the irreversibility of this process by positing that in an isolated system such as the gas entropy cannot decrease. The aim of statistical mechanics (SM) is to explain the behaviour of the gas, and in particular its conformity with the Second Law, in terms of the dynamical laws governing the individual molecules of which the gas is made up. In what follows these laws are assumed to be the ones of Hamiltonian classical mechanics.

*Department of Philosophy, Logic and Scientific Method, London School of Economics, Houghton Street, London WC2A 2AE, England, r.p.frigg@lse.ac.uk.

We should not, however, ask for an explanation of the the Second Law literally construed. This law is a universal law and as such cannot be explained by a statistical theory. But this is not a problem because we can rest content if we explain the ‘Boltzmannian version’ of the second law (Callender 1999), which I call ‘Boltzmann’s Law’ (BL):

Consider an arbitrary instant of time t and assume that at that time the Boltzmann entropy $S_B(t)$ of the system is low. It is then *highly probable* that at any time $t' > t$ we have $S_B(t') > S_B(t)$.

What notion of probability is invoked in BL and what reasons do we have to believe that the claim it makes is true? The orthodox answer is that probabilities are time averages and that entropy is likely to increase because, assuming that the system is ergodic, the system is in equilibrium most of the time. This view is now widely believed to be untenable due to both conceptual problems and its invocation of ergodicity (see Earman and Rédei (1996) and van Lith (2001) for discussions). A propensity interpretation of SM probabilities is ruled out by the fact that the underlying micro-theory, Hamiltonian mechanics, is deterministic: Popper’s and Miller’s hand-waving notwithstanding, this is incompatible with there being propensities (Clark 2001). Frequentism has never been seriously put forward as an interpretation of SM probabilities because mechanical systems do not satisfy von Mises’ independence requirement. Finally, so-called ‘no-theory-theories’ do not improve the situation because at least in the context of physical theories they do not provide an independent alternative to other accounts (Frigg and Hoefer 2007).

In two recent papers Loewer (2001, 2004) has suggested that the way out of this deadlock is to build on David Lewis’ (1986, 1994) approach and interpret SM probabilities as Humean chances. In this paper I first give a precise formulation of Loewer’s proposal, then raise two fundamental objections, and finally conclude that these can be overcome only at the price of interpreting SM probabilities epistemically.

2 Boltzmannian Statistical Mechanics

The microstate of a system consisting of n particles is specified by a point in its $6n$ -dimensional phase space Γ , which is endowed with the

Lebesgue measure μ_L (its ‘natural’ measure).¹ The dynamics of the system is governed by Hamilton’s equations of motion, which define a measure-preserving phase flow ϕ_t on Γ ; that is, $\phi_t: \Gamma \rightarrow \Gamma$ is a one-to-one mapping for every real number t and $\mu_L(\phi_t(B)) = \mu_L(B)$ for every measurable set $B \subseteq \Gamma$. In what follows we assume that the relevant physical process begins at a particular instant t_0 and I adopt the convention that ‘ $\phi_t(x)$ ’ denotes the state of the system at time $t_0 + t$ if it was in state x at t_0 , and likewise for ‘ $\phi_t(B)$ ’. Similarly, ‘ $\phi_{-t}(x)$ ’ denotes the state at time t_0 which gets mapped onto x at time $t_0 + t$ under the dynamics of the system, and likewise for $\phi_{-t}(B)$. In a Hamiltonian system energy is conserved and hence the motion of the system is confined to the $6n - 1$ dimensional energy hypersurface Γ_E . The measure μ_L can be restricted to Γ_E , which induces a natural and invariant measure μ on Γ_E .

The macrostates M_k , $k = 1, \dots, m$, of the system (where $m < \infty$), characterised by the values of macroscopic parameters, are assumed to supervene on the system’s microstates. Therefore each M_k is associated with a region $\Gamma_{M_k} \subseteq \Gamma_E$ so that the system is in macrostate M_k at t iff its microstate x at t lies within Γ_{M_k} . The Γ_{M_k} form a partition of Γ_E , meaning that they do not overlap and jointly cover Γ_E .

By definition, $S_B(M_k) := k_B \log[\mu(\Gamma_{M_k})]$ is the Boltzmann entropy of macrostate M_k (where k_B is the Boltzmann constant). Because the Γ_{M_k} don’t overlap it follows that a system is in exactly one macrostate at any given time t and for this reason it makes sense to talk about the Boltzmann entropy $S_B(t)$ of a system at time t : $S_B(t) := k_B \log[\mu(\Gamma_{M_t})]$, where M_t is the system’s macrostate at time t (i.e. M_t is the M_k for which it is the case that $x \in \Gamma_{M_k}$, where x is the system’s microstate at t). The Boltzmann entropy assumes its maximum for the equilibrium state.

To rationalise Boltzmann’s law we need to introduce probabilities. The standard way to do this is by appeal to the so-called ‘Statistical Postulate’ (SP):

Let M_t be the system’s macrostate at time t . Then the probability at time t that the system’s microstate lies in $B \subseteq \Gamma_{M_t}$ is $p_t(B) = \mu(B)/\mu(\Gamma_{M_t})$.

Now consider the set F of all microstates in Γ_{M_t} which in the near future evolve towards macrostates M' that have higher Boltzmann

¹For a short introduction to Boltzmannian SM see Lebowitz (1993).

entropy than M_t . With the assumption that $\mu(F)/\mu(\Gamma_{M_k}) \approx 1$ for all k it follows from SP that for all t the system is highly likely to evolve towards a state of higher entropy, which is exactly what BL asserts.

Whether or not this assumption holds true in a particular system is a substantial question. However, even if it does there is a problem. It follows from the time reversal invariance of Hamilton's equations of motion that if it is true that the system is overwhelmingly likely to evolve towards a macrostate of higher entropy in the future, it is also overwhelmingly likely to have evolved into the current macrostate *from* a past macrostate M'' which also has *higher* entropy. This flies in the face everyday experience and leads to wrong retrodictions.

Albert (2000, 71-96) suggests fixing this problem by first taking the system under investigation to be the entire universe and then adopting the so-called Past Hypothesis (PH), the postulate that the universe came into being in a low entropy macrostate, the Past State, which is provided to us by modern Big Bang cosmology. The problems with flawed retrodictions can then be avoided by conditionalising on the Past State. From a technical point of view, this amounts to replacing SP with what I call the 'Past Hypothesis Statistical Postulate' (PHSP):

Let M_t be the system's macrostate at time t . SP is valid for the Past State M_p , which obtains at time t_0 . For all times $t > t_0$ the probability at time t that the system's microstate lies in B is $p_t(B|R_t) = \mu(B \cap R_t)/\mu(R_t)$, where $R_t := M_t \cap \phi_t(M_p)$.

In what follows I refer to these probabilities as 'PHSP probabilities'. This principle is used to make predictions about the system's future by choosing B to be the set of those microstates that behave in the desired way. For instance, if choose B to be the set F (as defined above), $p_t(F)$ is the probability that the system's entropy will increase in the near future given PH. If this probability comes out high for all M_k we have explained BL. Again, whether or not this is the case is a substantive question having to do with both the construction of the macrostates as well as the dynamics of the system. However, the problems I discuss in this paper are orthogonal to this issue and so I assume for the sake of argument that this assumption bears out in systems of interest.²

²I also assume that one can make sense of PH in the current context, an assumption

3 Humean Chance

The basis for Lewis' theory of probability is the so-called Humean mosaic, the collection of all non-modal and non-probabilistic actual events making up the world's entire history (from the very beginning to the very end) and upon which all other facts supervene. Lewis himself suggested that the mosaic consists of space-time points plus local field quantities representing material stuff. In a classical mechanical system the Humean mosaic simply consist of the trajectory of the system's microstate in phase space, on which the system's macrostates supervene.

The next element of Lewis' theory is a thought experiment. To make this explicit – more explicit than it is in Lewis' own presentation – I introduce a fictitious creature, Lewis' Demon. In contrast to human beings who can only know a small part of the Humean mosaic, Lewis' Demon knows the entire mosaic. The demon now formulates various deductive systems which make true assertions about what is the case, and, perhaps, also about what the probability for certain events are. Then the demon is asked to choose the best among these systems. The laws of nature are the true theorems of this system and the chances for certain events to occur are what the probabilistic laws of the best system say they are (Lewis 1994, 480). Following Loewer, I call probabilities thus defined *L*-chances.

The best system is the one that strikes the best balance between strength, simplicity and fit. The notions of strength and simplicity are given to the demon and are taken for granted in this context, but the notion of fit needs explicit definition. Every system assigns probabilities to certain courses of history, among them the actual course; the fit of the system is measured by the probability that it assigns to the actual course of history, i.e. by how likely it regards things that actually happen. By definition systems that do not involve probabilistic laws have perfect fit. As an illustration, consider a Humean mosaic that consists of just ten outcomes of a coin flip: HHTHTTHHHT. Theory T_1 posits that all events are independent and sets $p(H) = p(T) = 0.5$; theory T_2 shares the independence assumption but posits $p(H) = 0.9$ and $p(T) = 0.1$. It follows that T_1 has better fit than T_2 because $(0.5)^{10} > (0.1)^5(0.9)^5$.

Loewer's suggestion is that Boltzmannian SM as introduced above

that has been questioned by Earman (2006).

– the package of Hamiltonian mechanics, PH and PHSP – is a putative best system of the sort just described (2001, 618; 2004, 1124) and that PHSP probabilities can therefore be regarded as Humean chances. But there is an obvious problem, namely reconciling determinism and the existence of probabilistic laws, which Lewis himself thought was impossible (1986, 118).

Loewer claims that Lewis was wrong about this and suggests that introducing probabilities via initial conditions solves the problem:

‘[...] while there are chances different from 0 and 1 for possible initial conditions the chances of any event A after the initial time will be either 1 or 0 since A ’s occurrence or non-occurrence will be entailed by the initial state and the deterministic laws. However, we can define a kind of dynamical chance which I call ‘macroscopic chance’. The macroscopic chance at t of event A is the probability given by starting with the micro-canonical distribution over the initial conditions and then conditionalising on the entire macroscopic history of the world (including the low entropy postulate) up until t . [...] this probability distribution is completely compatible with deterministic laws since it concerns only the initial conditions of the universe.’ (Loewer 2001, 618-19)³

Loewer does not tell us what exactly he means by ‘a kind of dynamical chance’, in what sense this chance is macroscopic, how its values are calculated, and how it connects to the technical apparatus of SM. I will now present how I think this proposal is best understood and show that, on this reading, Loewer’s ‘macroscopic chances’ coincide with PHSP as formulated above.

As in Section 2, I take the system’s state at $t > t_0$ to be the macrostate M_t . We now need to determine the probability of the event ‘being in set $B \subseteq \Gamma_{M_t}$ at time t ’. As I understand it, Loewer’s proposal falls into two parts. The first is that the probability of an event at a time t is ‘completely determined’ by the probability of the corresponding event at time t_0 ; that is, the probability of the event ‘being in set B at time t ’, $p_t(B)$, is equal to the probability of ‘being in set B_0 at time t_0 ’ where B_0 is, by definition, the set

³The same idea is described in Loewer (2004, 1124).

that evolves into B under the dynamics of the system after time t has elapsed. Formally, $p_t(B) = \mu_0(B_0) = \mu_0(\phi_{-t}(B))$, where μ_0 is the microcanonical distributino over the Past State, i.e. $\mu_0(\cdot) = \mu(\cdot \cap \Gamma_{M_p})/\mu(\Gamma_{M_p})$

The second part is conditionalising on the entire macro history up to time t , i.e a specification of the system's macro state at *each* instant of time between t_0 and t . A possible macro history, for instance, is that system is in macrostate M_1 during the interval $[t_0, t_1]$, in M_5 during $(t_1, t_2]$, in M_7 during $(t_2, t_3]$, etc., where t_1, t_2, t_3, \dots are the instants of time at which the system changes from one macrostate into another. What we are now expected to calculate is the probability of 'being in set B at time t ' *given* the system's macro history. Let Q_t be the set of all microstates in Γ_{M_t} that are compatible with the entire past history of the system; i.e. it is the set of all $x \in \Gamma_{M_t}$ that lie on trajectories that for every t were in the Γ_{M_k} corresponding to the *actual* macrostate of the system at t . The sought-after conditional probability then is $p_t(B|Q_t) = p_t(B \& Q_t)/p_t(Q_t)$, provided that $p_t(Q_t) \neq 0$, which, as we shall see, is the problematic condition.

Putting these two parts together we obtain the fundamental equation defining L -chances for deterministic systems:

$$p_t(B|Q_t) = \frac{\mu_0(\phi_{-t}(B \cap Q_t))}{\mu_0(\phi_{-t}(Q_t))}, \quad (1)$$

where, again, $\mu_0(\cdot) = \mu(\cdot \cap \Gamma_{M_p})/\mu(\Gamma_{M_p})$.

The crucial thing to realise now is that due to the conservation of the measure the expression for the conditional probability in PHSP can be expressed as $p_t(B|R_t) = \mu(\phi_{-t}(B \cap R_t))/\mu(\phi_{-t}(R_t))$. Trivially, we can substitute μ_0 for μ in this expression which makes it is equivalent to Equation 1, if we treat Q_t and R_t as equals. (In fact there is a difference between them in that R_t only involves a conditionalisation on PH, while Q_t contains the *entire* past history. However, nothing in PHSP depends on this and one could just as well include the entire history in R_t .) Hence PHSP can be interpreted as attributing probabilities to events at $t > t_0$ solely on the basis of the microcanonical measure over the initial conditions, which is precisely what Loewer needs.

4 Problems with Fit

Loewer claims that SM as introduced above is the system that strikes the best balance between simplicity, strength and fit. Trivially, this implies that it can be ranked along these three dimensions. Simplicity and strength are no more problematic in SM than they are in any other context and I shall therefore not discuss them further here. The problematic concept is fit.

The fit of a theory is measured in terms of the probability that it assigns to the actual course of history. But what history? Given that L -chances are calculated using the Lebesgue measure, which assigns measure zero to *any* trajectory, they do not lead to a non-trivial ranking of micro histories (trajectories in Γ). The right choice seems to be to judge the fit of theory with respect to the system's macro history.

What is the probability of a macro history? A first answer to this question would be to simply use Equation 1 to calculate the probability of a macro state at each instant of time and then multiply them all, just as we did in the above example with the coins (with the only difference that the probabilities are now not independent any more, which is accounted for in Equation 1). This is plain nonsense. There is an uncountable infinity of such probabilities and multiplying an uncountable infinity of numbers is an ill-defined operation. Determining the probability of a history by multiplying probabilities for individual events in the history works fine as long as the events are discrete (like coin flips), but it fails when we have a continuum.

Maybe this was too crude a stab at the problem and when taking the right sorts of limits things work out fine. Let us discretise time by dividing the real axis into small intervals of length δ , then calculating the probabilities at the instants $t_0, t_0 + \delta, t_0 + 2\delta$ etc., multiply them (there are only countably many now), and then take the limit $\delta \rightarrow 0$. This would work if the $p_t(B|Q_t)$ depended in a way on δ that would assure that the limit exists. This is not the case. In fact, for all $t > t_1$ (i.e. after the first change of macrostate), the $p_t(B|Q_t)$ do not exist because Q_t has measure zero, and this irrespective of δ . This can be seen as follows. Take the above example of a macro history and consider an instant $t \in (t_1, t_2]$ when the system is in macrostate M_5 . To calculate the probability of the system being in M_5 at t we need to determine Q_t , the set of all microstates in Γ_{M_5} compatible with the past macro history. Now, these points must be such that they were

were in M_1 at t_1 and in M_5 just an instant later (i.e. for any $\varepsilon > 0$, at $t_1 + \varepsilon$ the system's state is in Γ_{M_5}). The mechanical systems we are considering have global solutions (or at least solutions for the entire interval $[t_0, t_f]$, where t_f is the time when the system ceases to exist) and trajectories in such systems have finite phase velocity; that is, a phase point x in Γ cannot cross a finite distance in no time. From this it follows that the only points that satisfy the condition of being in M_1 at t_1 and in M_5 just instant later are the ones that at t_1 lie exactly on the boundary between M_1 and M_5 . But the boundary of a $6n - 1$ dimensional region is $6n - 2$ dimensional and therefore has measure zero. Therefore Q_t has measure zero for all $t > t_1$, and accordingly $p_t(B|Q_t)$ does not exist for $t > t_1$, no matter what B is. Needless to say, this renders the limit $\delta \rightarrow 0$ obsolete.

The source of the problem is the conjunction of three elements: (1) the posit that time is continuous, (2) the assumption that the transition from one macrostate to another one takes place at a precise instant, (3) the posit that we conditionalise on the *entire* macro history of the system. We have to give up at least one of these to obtain non-zero $p_t(B|Q_t)$. The problem is that all three elements either seem reasonable or are deeply entrenched in the theory and cannot be renounced without far-reaching consequences.

The first option, discretising time, would solve the problem because if we assume that time is discrete the macro history is discrete too. If we only consider, say, the events 'being in Γ_{M_1} at instant τ_1 ' and 'being in Γ_{M_5} at instant τ_2 ', where $\tau_1 \in [t_0, t_1]$ and $\tau_2 \in (t_1, t_2]$, sets of finite measure can move from Γ_{M_1} to Γ_{M_5} and Q_t no longer needs to have measure zero. The problem with this suggestion is that it is ad hoc and defeats the purpose of SM. If we believe that classical mechanics is the fundamental theory governing the micro constituents of the universe and set out to explain the behaviour of the universe in terms of its laws, not much seems to be gained if such an explanation can only be had at the expense of profoundly modifying these laws.

The second suggestion would be to allow for finite transition times between macrostates, i.e. allowing for there to be periods during which it is indeterminate in which macrostate the system is. Given that states move with finite phase velocity, this would amount to introducing 'transition zones', i.e. 'belts' between different Γ_{M_k} consisting of microstates which neither belong to one nor the other macrostate. This suggestion is not without merit as one could argue that sharp

boundaries between macro-states are indeed a mathematical idealisation that is ultimately unjustifiable from a physics perspective. However, giving up the assumption that the Γ_{M_k} have sharp boundaries and together form a partition of Γ_E would amount to a serious change in SM itself, and it would remain to be seen whether a theory based on this assumption turns out to be workable.

The third option denies that we should conditionalise on the *complete* macro history. The idea is that even though time at bottom is continuous, the macro history takes record of the system's macrostate only at discrete instants and is oblivious about what happens between these. That is, what we should conditionalise on is the discrete macro history (DMH): M_p at τ_0 , M_{τ_1} at τ_1 , M_{τ_2} at τ_2 , ..., $M_{\tau_{j-1}}$ at τ_{j-1} and M_{τ_j} at τ_j , where $t_0 =: \tau_0 \leq \tau_1 \leq \dots \leq \tau_{j-1} \leq \tau_j := t_f$ is a finite number of instants of time and the M_{τ_i} the system's macro state at time τ_i .⁴ This solves the problem because when conditionalising on a discrete macro history Q_t no longer necessarily is of measure zero.

This is at once the most feasible and the most problematic suggestion. It is feasible because it does not require revisions in the structure of the theory. It is problematic because we have given up the notion that the fit of a theory has to be best with respect to the complete history of the world, and replaced it with the weaker requirement that fit be best for a partial history. (And mind you, the point is not that the fit of the full history is in practice too complicated to calculate and we therefore settle for a more tractable notion; the point is that the fit of a complete macro history is simply not defined because the relevant conditional probabilities do not exist.) From the point of view of Lewis' theory this seems unmotivated. Fit, like truth, is a semantic concept characterising the relation between the theory and the world, and if the Humean mosaic has continuous events in it there should still be a matter of fact about what the fit of the theory is.

Moreover, even if one was willing to believe that a discrete version of fit was satisfactory, it is not clear whether this leads to useful results. Depending on which particular instants of time one chooses to measure fit, one can get widely different results. Conditionalising on a DMH would be useful only if it was the case that the fit rankings came out the same no matter what choice of instants we make. There is at least

⁴I assume j to be finite. There is a further problem with infinite sequences (Elga 2004). The difficulties I discuss in this section and the next are independent of that problem and Elga's solution is available also in the present context.

a question whether this is the case.

5 The Putative Best System Is Not the Best System

I now assume that a DMH notion of fit can be defended in one way or another.⁵ Then a further problem emerges: the package consisting of Hamiltonian Mechanics, PH and PHSP in fact is *not* the best system. The reason for this is that we can always improve the fit of a system if we choose a measure that, rather than being uniform over Γ_{M_p} , is somehow peaked over those initial conditions that are compatible with the entire DMH.

Let us make this more precise. The probability of the DMH is $p(DMH) = p_{\tau_0}(B_{\tau_0}|Q_{\tau_0}) \dots p_{\tau_{j-1}}(B_{\tau_{j-1}}|Q_{\tau_{j-1}})$, where the B_{τ_i} are those subsets of M_{τ_i} that evolve into $M_{\tau_{i+1}}$ under the evolution of the system. One can then prove that

$$p(DMH) = \mu_0[\Gamma_p \cap \phi_{-\tau_1}(\Gamma_1) \cap \dots \cap \phi_{-\tau_j}(\Gamma_j)] \quad (2)$$

where $\Gamma_p := \Gamma_{M_p}$ and $\Gamma_i := \Gamma_{M_{\tau_i}}$ for $i = 1, \dots, j$. Now define $N := \Gamma_p \cap \phi_{-\tau_1}(\Gamma_1) \dots \phi_{-\tau_j}(\Gamma_j)$. The fit of system is measured by the probability that it assigns to the actual DMH, which is given by Equation 2. It is a straightforward consequence of this equation that the fit of a system can be improved by replacing μ_0 by a measure μ_P that is peaked over N , i.e. $\mu_P(N) > \mu_0(N)$ and $\mu_P(\Gamma_p \setminus N) < \mu_0(\Gamma_p \setminus N)$ while $\mu_P(\Gamma_p) = \mu_0(\Gamma_p)$. Fit becomes maximal (i.e. $p(DMH) = 1$) if, for instance, we choose the measure μ_N that assigns all the weight to N and none to $\Gamma_p \setminus N$: for any set $B \subseteq \Gamma_p$ we have $\mu_N(B) := k\mu_0(B \cap N)$, where $k = 1/\mu_0(N)$ (provided that $\mu_0(N) \neq 0$). Trivially, N contains the *actual* initial condition of the universe. A simpler and more convenient distributions that yields maximal fit is a Dirac delta function over the actual initial condition.

If there is such a simple way to improve (and even maximise) fit, why does the demon not provide us with a system comprising μ_N or a delta function? Coming up with such a system is not a problem for

⁵I make this choice for convenience; the problem that I describe in this section also arises for the other two options.

the demon, as, by assumption, he knows the entire Humean mosaic, which contains the exact initial condition.

A reason to prefer μ_0 to other measures might be that these make the system less simple and that this loss in simplicity is not compensated by a corresponding gain in fit and strength. This seems implausible. Handling a dirac delta function rather than μ_0 does not render the system more complicated while the gain in fit is considerable. Hence simplicity does not seem to provide reason to prefer μ_0 to other measures that have better fit.

6 Outlook: Epistemic Probabilities After All

The system consisting of Hamiltonian mechanics, PH, and PHSP is not the best system and therefore PHSP probabilities cannot be interpreted as Humean chances. In this section I first want suggest that the probabilities in this system are best understood as epistemic probabilities of sorts and then indicate how this view could be defended against some common objections.

Every theory involving probabilities must answer the question of what these probabilities are probabilities for. The initial conditions approach to chance does not seem to have an answer to this question. The universe has exactly one initial micro condition and there is nothing chancy about this condition. How, then, can we understand a probability distribution over initial conditions? The only answer seems to be that this distribution reflects our ignorance about the systems's initial micro condition; all we know is the system's initial macrostate and so we put a probability distribution over the micro conditions compatible with that macrostate which reflects our lack of knowledge.

How these epistemic probabilities should be understood is a question that I cannot discuss here. Let me just indicate that there are at least two options. The first, a version of objective Bayesianism, appeals to Jayenes' maximum entropy principle, which indeed instructs us to prefer μ_0 to alternative measures because, given the information about the system's macro state, μ_0 maximises the (continuous) Shannon entropy. The other alternative is to revise Lewis' account in a way that builds epistemic restrictions of the users of theories into

the selection criteria for systems. Hoefer's (2007) theory of Humean chance makes room for this possibility.

There are two main complaints about an epistemic interpretation of SM probabilities. The first points out that the thermodynamic entropy is a property of a physical system and that S_B coincides with it up to a constant. This, so the argument goes, is inexplicable on the basis of an epistemic approach to probabilities.⁶ This is wrong because S_B is defined in terms of the measure of certain chunks of phase space and probabilities (no matter how we interpret them) have simply nothing to do with it.

The second complaint concerns the alleged causal efficacy of human knowledge. The point becomes clear in the following – rhetorical – questions by Albert:⁷

‘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 into*? 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, 64, original emphasis)

It can't be, and no one should think that it could. Proponents of epistemic probabilities need not believe in parapsychology. What underlies this objection is the mistaken view that PHSP probabilities play a part in bringing about things in the world. Of course the cooling down of drinks and the boiling of kettles has nothing to do with what anybody thinks or knows about them; but they have nothing to do with the probabilities attached to these events either. Drinks cool down and kettles boil because the universe's initial condition is such that under the dynamics of the system it evolves into a state in which this happens. All we need to explain why things happen is the initial condition and the dynamics.

Last but not least, the decision to conditionalise on DMH rather than the full macro history seems to square better with an epistemic

⁶This point is often made in conversation, but I have been unable to locate it in print.

⁷Redhead (1995, 27-28, 32-33), Loewer (2001, 611), Goldstein (2001, 48), and Meacham (2005 287-8) make similar points.

approach to probabilities. For these reasons I suggest that we take seriously the option to interpret SM probabilities epistemically.

Acknowledgements

Thanks to Nancy Cartwright, David Lavis, Carl Hoefer, Barry Loewer, and Charlotte Werndl for helpful discussion and/or comments on earlier drafts. Thanks also to the audiences of the conference ‘Time, Chance and Reduction’ in Munich and the PSA Meeting in Vancouver for feedback and suggestions.

Bibliography

- Albert, David (2000), *Time and Chance*. Cambridge/MA and London: Harvard University Press.
- Callender, Craig (1999), “Reducing Thermodynamics to Statistical Mechanics: The Case of Entropy”, *Journal of Philosophy* 96: 348-373.
- Earman, John (2006), “The ‘Past Hypothesis’: Not Even False”, *Studies in History and Philosophy of Modern Physics* 37: 399-430.
- and Miklós Rédei (1996), “Why Ergodic Theory Does Not Explain the Success of Equilibrium Statistical Mechanics”, *British Journal for the Philosophy of Science* 47: 63-78.
- Elga, Adam (2004), “Infinitesimal Chances and the Laws of Nature”, *Australasian Journal of Philosophy* 82: 67-76.
- Frigg, Roman and Carl Hoefer (2007), “Probability in GRW Theory”, *Studies in History and Philosophy of Science* 38: 371-389.
- Goldstein, Sheldon (2001), “Boltzmann’s Approach to Statistical Mechanics”, in: Bricmont et al. (2001): *Chance in Physics: Foundations and Perspectives*. Berlin and New York: Springer, 39-54.
- Hoefer, Carl (2007), “The Third Way on Objective Probability: A Sceptic’s Guide to Objective Chance”, *Mind* 116: 549-596.
- Lebowitz, Joel (1993), “Macroscopic Laws, Microscopic Dynamics, Time’s Arrow and Boltzmann’s Entropy”, *Physica A* 194: 1-27.

- Lewis, David (1986), “A Subjectivist’s Guide to Objective Chance” and ‘Postscripts to “A subjectivist’s guide to objective Chance”’, in David Lewis, *Philosophical Papers*, Vol. 2, Oxford: Oxford University Press, 83-132.
- (1994), “Humean Supervenience Debugged”, *Mind* 103: 473-90.
- Loewer, Barry (2001): “Determinism and Chance”, *Studies in History and Philosophy of Modern Physics* 32: 609-629.
- (2004), “David Lewis’ Humean Theory of Objective Chance”, *Philosophy of Science* 71: 1115 -1125.
- Meacham, Christopher (2005), “Three Proposals Regarding a Theory of Chance”, *Philosophical Perspectives* 19: 281-307.
- Redhead, Michael (1995), *From Physics to Metaphysics*. Cambridge: Cambridge University Press.
- van Lith, Janneke (2001), “Ergodic Theory, Interpretations of Probability and the Foundations of Statistical Mechanics”, *Studies in History and Philosophy of Modern Physics* 32: 581-594.