

6

Probability in Boltzmannian statistical mechanics

ROMAN FRIGG

In: Gerhard Ernst and Andreas Hüttemann (eds.): *Time, Chance and Reduction. Philosophical Aspects of Statistical Mechanics*. Cambridge: Cambridge University Press, 2010, 92-118.

6.1 Introduction

A cup of tea, left to its own, cools down while the surrounding air heats up until both have reached the same temperature, and a gas, confined to the left half of a room, uniformly spreads over the entire available space as soon as the confining wall is removed. Thermodynamics (TD) characterizes such processes in terms of an increase of thermodynamic entropy, which attains its maximum value at equilibrium, and the second law of thermodynamics posits that in an isolated system entropy cannot decrease. The aim of statistical mechanics (SM) is to explain the behaviour of these systems, in particular their conformity with the second law, in terms of the dynamical laws governing the individual molecules of which the systems are made up. In what follows these laws are assumed to be the ones of classical mechanics.

An influential suggestion of how this could be achieved was made by Ludwig Boltzmann (1877), and variants of it are currently regarded by many as the most promising option among the plethora of approaches to SM. Although these variants share a commitment to Boltzmann's basic ideas, they differ widely in how these ideas are implemented and used. These differences become most tangible when we look at how the different approaches deal with probabilities. There are two fundamentally different ways of introducing probabilities into SM, and even within these two groups there are important disparities as regards both technical and interpretational issues. The aim of this chapter is to give a statement of these approaches and point out wherein their difficulties lie.

6.2 Boltzmannian statistical mechanics

The state of a classical system of n particles with three degrees of freedom is completely specified by a point $x = (\mathbf{r}_1, \mathbf{p}_1, \dots, \mathbf{r}_n, \mathbf{p}_n)$ in its $6n$ dimensional phase space Γ , where \mathbf{r}_i and \mathbf{p}_i are position and momentum, respectively, of the i th

particle.¹ For reasons that will become clear soon, x is referred to as the system's (fine-grained) micro-state. Γ is endowed with the 'standard' Lebesgue measure μ . Because we only consider systems in which the total energy is conserved and which have finite spatial extension, only a finite region $\Gamma_a \subset \Gamma$ is accessible to x .

We now assume that dynamics of the system is governed by Hamilton's equations, which define a measure-preserving flow ϕ_t on Γ : for all times t , $\phi_t : \Gamma \rightarrow \Gamma$ is a one-to-one mapping and $\mu(R) = \mu(\phi_t(R))$ for all regions $R \subseteq \Gamma$.² To indicate that we consider the image of a point x under the dynamics of the system between two instants of time t_1 and t_2 (where $t_1 < t_2$) I write ' $\phi_{t_2-t_1}(x)$ ', and likewise for ' $\phi_{t_2-t_1}(R)$ '. The inverse is denoted by ' $\phi_{t_1-t_2}(x)$ ' and provides the system's state at time t_1 if its state was x at t_2 , and likewise for ' $\phi_{t_1-t_2}(R)$ '.

It is one of the central assumptions of Boltzmannian SM (BSM) that the system's macro-states M_i , $i = 1, \dots, m$ (and $m < \infty$) – characterized by the values of macroscopic variables such as pressure, volume, and temperature – supervene on the system's micro-states, meaning that a change in the macro-state must be accompanied by a change in the micro-state. Hence the system's micro-state uniquely determines the system's macro-state in that to every given $x \in \Gamma_a$ there corresponds *exactly one* macro-state, $M(x)$. This determination relation is not one-to-one as many different x can correspond to the same macro-state. It is therefore natural to define the macro-region of M_i , Γ_{M_i} , as the subset of Γ_a that consists of all micro-states corresponding to macro-state M_i :

$$\Gamma_{M_i} := \{x \in \Gamma_a \mid M_i = M(x)\}, i = 1, \dots, m. \quad (6.1)$$

The Γ_{M_i} together form a partition of Γ_a , meaning that they do not overlap and jointly cover Γ_a : $\Gamma_{M_i} \cap \Gamma_{M_j} = \emptyset$ for all $i \neq j$ and $i, j = 1, \dots, m$, and $\Gamma_{M_1} \cup \dots \cup \Gamma_{M_m} = \Gamma_a$, where ' \cup ', ' \cap ' and ' \emptyset ' denote set theoretic union, intersection and the empty set respectively.

The Boltzmann entropy of a macro-state M is defined as

$$S_B(M) = k_B \log[\mu(\Gamma_M)], \quad (6.2)$$

where M ranges over the M_i and k_B is the so-called Boltzmann constant. The macro state for which S_B is maximal is the equilibrium state, meaning that the system is

¹ For an introduction to classical mechanics see, for instance, Goldstein (1981) and Abraham and Marsden (1980); Goldstein (2001), Goldstein and Lebowitz (2004) and Lebowitz (1993a, 1993b, 1999) provide compact statements of the Boltzmannian approach to SM.

² In a Hamiltonian system energy E is conserved and hence the motion of the system is confined to the $6n-1$ dimensional energy hypersurface Γ_E defined by the condition $H(x) = E$, where $H(x)$ is the Hamiltonian of the system. Restricting μ to Γ_E yields a natural measure μ_E on this hypersurface. At some points in the argument below it would be more appropriate to use μ_E rather than μ . However, in the specific circumstances this would need some explaining and as none of the conclusions I reach depend on this I will not belabour this point further. For further discussion of this point, see Frigg (2008: 103–113).

in equilibrium if it has reached this state.³ For the sake of notational convenience we denote the equilibrium state by M_{eq} and choose, without loss of generality, the labelling of the macro-states such that $M_m = M_{eq}$.

Given this, we can define the Boltzmann entropy of a *system* at time t as the entropy of the system's macro-state at t :

$$S_B(t) := S_B(M_t) \tag{6.3}$$

where M_t is the system's macro-state at time t ; i.e. $M_t := M(x(t))$, where $x(t)$ is the system's micro-state at t .

It is now common to accept the so-called past hypothesis, the postulate that the system starts off at t_0 in a low-entropy macro-condition, the 'past state' M_p (and we choose our labelling of macro-states such that $M_1 = M_p$). How the past state is understood depends on one's views on the scope of SM. The grand majority of Boltzmannians take the system under investigation to be the entire universe and hence interpret the past state as the state of the universe just after the big bang (I will come back to this below). Those who stick to the spirit of laboratory physics regard laboratory systems as the relevant unit of analysis and see the past state as one that is brought about in some particular experimental situation (such as the gas being confined to the left half of the room). For the discussion in this chapter it is inconsequential which of these views is adopted.

The leading idea of BSM is that the behaviour of $S_B(t)$ should mirror the behaviour of the thermodynamic entropy S_{TD} ; that is, it should increase with time t and reach its maximum at equilibrium. We should not, however, expect that this mirroring be exact. The second law of thermodynamics is a universal law and says that S_{TD} can never decrease. A statistical theory cannot explain such a law and we have to rest content if we explain the 'Boltzmannian version' of the second law (Callender, 1999), also referred to as '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 $S_B(t'') \geq S_B(t')$.

A system that behaves in accordance with BL is said to exhibit 'thermodynamic like behaviour' (TD-like behaviour, for short).

What notion of probability is invoked in BL and what reasons do we have to believe that the claim it makes is true? Different approaches to BSM diverge in how they introduce probabilities into the theory and in how they explain the tendency

³ This assumption is problematic for two reasons. First, Lavis (2005) has argued that associating equilibrium with *one* macro-state is problematic for different reasons. Second, that the equilibrium state is the macro-state with the highest Boltzmann entropy is true only for non-interacting systems like ideal gases (Uffink, 2007: section 4.4). I set these difficulties aside for now.

of $S_B(t)$ to increase. The most fundamental distinction is between approaches that assign probabilities directly to the system's macro-states, and approaches that assign probabilities to the system's micro-state being in a particular subset of the macro-region corresponding to the system's current macro-state; for want of better terms I refer to these as 'macro-probabilities' and 'micro-probabilities' respectively. I now present these approaches one at a time and examine whether, and if so to what extent, they succeed in explaining BL.

Before delving into the discussion, let me list those approaches to probability that can be discounted straight away within the context BSM, irrespective of how exactly probabilities are introduced into the theory. The first three items on this list are the classical interpretation, the logical interpretation and the so-called no-theory theory. These have not been put forward as interpretations of Boltzmannian probabilities, and this for good reasons. The first two simply are not the right kind of theories, while the no-theory does not seem to offer a substantial alternative to either the propensity theory or David Lewis' view (Frigg and Hoefer, 2007: 385, 386). Frequentism, as von Mises himself pointed out, is problematic as an interpretation of SM probabilities because a sequence of results that is produced by the same system does not satisfy the independence requirements of a collective (van Lith, 2001: 587). Finally, a propensity interpretation is ruled out by the fact that the underlying micro-theory, classical mechanics, is deterministic,⁴ which is incompatible with there being propensities (Clark, 2001).

6.3 The macro-probability approach

In this section I discuss Boltzmann's (1877) proposal to assign probabilities to the system's macro-states and Paul and Tatiana Ehrenfest's (1912) view that these should be interpreted as time averages, based on the assumption that the system is ergodic.

6.3.1 Introducing macro-probabilities

Boltzmann's way of introducing probabilities into SM is intimately related to his construction of the macro-regions Γ_{M_i} , which is now also known as the 'combinatorial argument'.⁵ Assume now that the system under investigation is a gas that consists of n molecules of the same type and is confined to a finite volume V . Then consider the six-dimensional phase space of *one* gas molecule, commonly referred

⁴ This is true of the systems studied in BSM; there are failures of determinism in classical mechanics when we look at a larger class of systems (Earman, 1986).

⁵ For a presentation and discussion of this argument see Uffink (2007: 974–893).

to as μ -space. The conservation of the total energy of the system and the fact that the gas is confined to V results in only a finite part of the μ -space being accessible to the particle's state. Now put a grid-like partition on μ -space whose cells all have the size $\delta\omega$ and whose borders run in the directions of the momentum and position axes. This results in the accessible region being partitioned into a finite number of cells ω_i , $i = 1, \dots, k$. The state of the entire gas can be represented by n points in μ -space, every one of which comes to lie within a particular cell ω_i . An 'arrangement' is a specification of which point lies in which cell; i.e. it is a list indicating, say, that the state of molecule 1 lies in ω_9 , the state of molecule 2 lies in ω_{13} , and so on. A 'distribution' is a specification of how many points (no matter which ones) are in each cell; i.e. it is a k -tuple (n_1, \dots, n_k) , expressing the fact that n_1 points are in cell ω_1 , and so on. Elementary combinatorial considerations show that each distribution is compatible with

$$W(n_1, \dots, n_k) := \frac{n!}{n_1! \dots n_k!} \quad (6.4)$$

arrangements. Boltzmann then regards the probability $p(n_1, \dots, n_k)$ of a distribution as determined by $W(n_1, \dots, n_k)$:

The probability of this distribution is then given by the number of permutations of which the elements of this distribution are capable, that is by the number $[W(n_1, \dots, n_k)]$. As the most probable distribution, i.e. as the one corresponding to thermal equilibrium, we again regard that distribution for which this expression is maximal [...].

*(Boltzmann, 1877: 187)*⁶

In other words, Boltzmann's posit is that $p(n_1, \dots, n_k)$ is proportional to $W(n_1, \dots, n_k)$.

Macro-states are determined by distributions because the values of a system's macroscopic variables are fixed by the distribution. That is, each distribution corresponds to a macro-state.⁷ To simplify notation, let us now assume that all distributions are labelled with an index i and that there are m of them. Then there corresponds a macro-state M_i to every distribution D_i , $i = 1, \dots, m$. This is sensible also from a formal point of view because each distribution corresponds to a particular region of Γ ; regions corresponding to different distribution do not overlap and all regions together cover the accessible parts of Γ (as we expect it to be

⁶ This and all subsequent quotes from Boltzmann are my own translations. Square brackets indicate that Boltzmann's notation has been replaced by the one used in this chapter. To be precise, in the passage here quoted 'W' refers to distribution over a partitioning of energy, not phase space. However, when Boltzmann discusses the partitioning of μ -space a few pages further down (on p. 191) he refers the reader back to his earlier discussions that occur on p. 187 (quoted) and p. 176. Hence he endorses the posits made in this quote also for a partitioning of phase space.

⁷ It may turn out that it is advantageous in certain situations to regard two or more distributions as belonging to the same macro-state (e.g. if the relevant macro-variables turn out to have the same values for all of them). As this would not alter any of the considerations to follow, I disregard this possibility henceforth.

the case for macro-regions, see Section 6.2). One can then show that the measure of each macro-region thus determined is given by

$$\mu(\Gamma_{M_i}) = W(n_1, \dots, n_k) (\delta\omega)^n \quad (6.5)$$

where $D_i = (n_1, \dots, n_k)$ is the distribution corresponding to M_i .

This allows us to restate Boltzmann's postulate about probabilities in terms of the measures of the macro-regions Γ_{M_i} :

The probability of macro-state M_i is given by

$$p(M_i) = c \mu(\Gamma_{M_i}), \quad i = 1, \dots, m \quad (6.6)$$

where c is a normalization constant determined by the condition $\sum_{i=1}^m p(M_i) = 1$.

I refer to this postulate as the 'proportionality postulate' and to probabilities thus defined as 'macro-probabilities'. The choice of the former label is evident; the latter is motivated by the fact that the postulate assigns probabilities to the macro-states of the system and that the values of these probabilities are determined by the measure of the corresponding macro-region. It is a consequence of this postulate that the equilibrium state is the most likely state.

Boltzmann does not provide an interpretation of macro-probabilities and we will return to this problem below. Let us first introduce his explanation of TD-like behaviour. The leading idea of Boltzmann's account of non-equilibrium SM is to explain the approach to equilibrium by a tendency of the system to evolve from an unlikely macro-state to a more likely macro-state and finally to the most likely macro-state:

In most cases the initial state will be a very unlikely state. From this state the system will steadily evolve towards more likely states until it has finally reached the most likely state, i.e. the state of thermal equilibrium.

(Boltzmann, 1877: 165)

[...] the system of bodies always evolves from a more unlikely to a more likely state.

(Boltzmann, 1877: 166)

In brief, Boltzmann's answer to the question of how SM explains the second law is that it lies in the nature of a system to move from states of lower towards states of higher probability.

6.3.2 Macro-probabilities scrutinized

The question now is where this tendency to move towards more probable states comes from. It does not follow from the probabilities themselves. The $p(M_i)$ as defined by Equation (6.6) are unconditional probabilities, and as such they do not imply anything about the succession of macro-states, let alone that ones of low

probability are followed by ones of higher probability. As an example consider a loaded die; the probability to get a ‘six’ is 0.25 and all other numbers of spots have probability 0.15. Can we then infer that after, say, a ‘three’ we have to get a ‘six’ because the six is the most likely event? Of course not; in fact, we are much more likely not to get a ‘six’. And the situation does not change if, to make the scenario more like SM, the dice is so strongly biased that getting a ‘six’ is by far the most likely event. If the probability for this to happen is 0.999, say, then we are of course very likely to see a ‘six’ when throwing the die next, but this has nothing to do with the fact that we threw a ‘three’ before. The expectation that we move through progressively more probable events on many rolls is simply a delusion.

A further (yet related) problem is that BL makes a statement about a *conditional* probability, namely the probability of the system’s macro-state at t'' , $M(t'')$, being such that $S_B(t'') > S_B(t')$, given that the system’s macro-state at an earlier time t' was such that its Boltzmann entropy was $S_B(t')$. The probabilities of the proportionality postulate are not of this kind. But maybe this mismatch is only apparent because conditional probabilities can be obtained from unconditional ones by using $p(B|A) = p(B \& A)/p(A)$, where A and B are arbitrary events and $p(A) > 0$. The relevant probabilities then would be $p(M_j|M_i) = p(M_j \& M_i)/p(M_i)$. Things are not that easy, unfortunately. M_j and M_i are mutually exclusive (the system can only be either in M_j or M_i but not in both) and hence the numerator of this relation is always zero.

The problem with this attempt is, of course, that it does not take time into account. What we really would have to calculate are probabilities of the form $p(M_j \text{ at } t'' \& M_i \text{ at } t')/p(M_i \text{ at } t')$. The problem is that it is not clear how to do this on the basis of the proportionality postulate as time does not appear in this postulate at all. One way around this problem might be to slightly revise the postulate by building time dependence into it: $p(M_i \text{ at } t) = c \mu(\Gamma_{M_i})$, $i = 1, \dots, m$. But even if this was justifiable, it would not fit the bill because it remains silent about how to calculate $p(M_j \text{ at } t'' \& M_i \text{ at } t')$. In sum, the probabilities provided to us by the proportionality postulate are of no avail in explaining BL.

This becomes also intuitively clear once we turn to the problem of interpreting macro-probabilities. Boltzmann himself remains silent about this question and the interpretation that has become the standard view goes back to the Ehrenfests (1912). They suggest to interpret the $p(M_i)$ as time averages, i.e. as the fraction of time that the system spends in Γ_{M_i} .⁸ The equilibrium state, then, is the most

⁸ Although time averages are, loosely speaking, the ‘continuum version’ of frequencies, there are considerable differences between the two interpretations. For a discussion of this point see von Plato (1988: 262–65; 1989: 434–37).

probable state in the sense that the system is in equilibrium most of the time. Two things are worth noticing about this interpretation. First, it makes it intuitively clear why macro-probabilities cannot explain TD-like behaviour: from the fact that the system spends a certain fraction of time in some macro-state M_i nothing follows about which state the system assumes after leaving M_i – time averages have no implications for the succession of states. Second, this interpretation needs to be justified by a particular dynamical explanation; whether the time the system spends in a macro-state is proportional to the volume of the corresponding macro-region, depends on the dynamics of the system.

Of what kind does that dynamics have to be for a time average interpretation to be possible? Again, Boltzmann himself remains silent about the dynamical conditions necessary to backup his take on probability.⁹ The Ehrenfests (1912) fill this gap by attributing to him the view that the system has to be ergodic. Roughly speaking, a system is ergodic on Γ_a if for almost all trajectories, the fraction of time a trajectory spends in a region $R \subseteq \Gamma_a$ equals the fraction of the area of Γ_a that is occupied by R .¹⁰ If the system is ergodic, it follows that the time that its actual micro-state spends in each Γ_{M_i} is proportional to $\mu(\Gamma_{M_i})$, which is what we need.

This proposal suffers from various technical difficulties having to do with the mathematical facts of ergodic theory. These problems are now widely known and need not be repeated here.¹¹ What needs to be noted here is that even if all these problems could be overcome, we still would not have explained BL because macro-probabilities, no matter how we interpret them, are simply the wrong probabilities to explain TD-like behaviour.¹²

Or would we? When justifying the time average interpretation we introduced the assumption that the system is ergodic. Ergodicity implies that a system that starts off in the past state sooner or later reaches equilibrium and stays in equilibrium most of the time. Is that not enough to rationalize BL? No it is not. What needs to be shown is not only that the system sooner or later reaches equilibrium, but also that this approach to equilibrium is such that whenever the system's macro-state changes the entropy is overwhelmingly likely to increase. To date, no one succeeded in showing that anything of that kind follows from ergodicity (or as Jos Uffink (2007: 981) puts the point, no proof of a statistical H -theorem on the basis of the ergodic hypothesis has been given yet). And there are doubts whether

⁹ And, as Uffink (2007: 981) points out, he reaffirmed later in 1881 that he did not wish to commit to any dynamical condition in his 1877 paper.

¹⁰ For a rigorous introduction to ergodic theory see Arnold and Avez (1968), and for an account of its long and tangled history Sklar (1993: Chs. 2 and 5) and von Plato (1994: Ch. 3).

¹¹ See Sklar (1993: Ch. 5), Earman and Rédei (1996) and van Lith (2001) for discussions.

¹² In line with the majority of Boltzmannians I here frame the problem of SM as explaining why entropy increases when a system is prepared in a low-entropy state. Lavis (2005: 254–61) criticizes this preoccupation with 'local' entropy increase as misplaced and suggests that what SM should aim to explain is that the Boltzmann entropy is close to its maximum most of the time.

such a proof is possible. Ergodicity is compatible with a wide range of different dynamical laws and it is in principle conceivable that there is a dynamical law that is such that the system passes through different macro-states in a way that is non-TD-like. To rule this out (and thereby justify BL on the basis of the ergodic hypothesis), one would have to show that the macro-state structure defined by the combinatorial argument is such that there exists no ergodic system that passes these macro-states in a non-TD-like way. It seems unlikely that there is an argument for this conclusion.

Even though ergodicity itself does not fit the bill, an important lesson can be learned from these considerations, namely that the key to understanding the approach to equilibrium lies in dynamical considerations. Assume for a minute that ergodicity was successful in justifying TD-like behaviour. Then, the explanation of TD-like behaviour would be entirely in terms of the dynamical properties of the system and the structure of the macro-regions; it would be an account of why, given ergodicity, the system visits the macro-states in the ‘right’ (i.e. TD-like) order. The proportionality principle and the probabilities it introduces would play no role in this; an explanation of the system’s behaviour could be given without mentioning probabilities once. Of course, ergodicity is not the right condition, but the same would be the case for any dynamical condition. What does the explaining is an appeal to features of the system’s phase flow in relation to the partitioning of the phase space into macro-regions Γ_{M_i} ; probabilities have simply become an idle wheel.

6.4 The micro-probability approach

David Albert (2000) proposed a different approach to probabilities in BSM. In this section I discuss this approach and Barry Loewer’s suggestion (2001; 2004) to understand these probabilities as Humean chances in David Lewis’ (1986; 1994) sense.

6.4.1 Introducing micro-probabilities

We now assign probabilities not to the system’s macro-states, but to measurable subsets of the macro-region corresponding to the system’s macro-state at time t . For lack of a better term I refer to these probabilities as ‘micro-probabilities’. An obvious way of introducing micro-probabilities is the so-called ‘statistical postulate’ (SP):

Let M_t be the system’s macro-state at time t . Then the probability at time t that the system’s micro-state lies in $A \subseteq \Gamma_{M_t}$ is $p_t(A) = \mu(A)/\mu(\Gamma_{M_t})$.

Given this, we can now calculate the probabilities occurring in BL. Let $\Gamma_{M_{t'}}^{(+)}$ be the subset $\Gamma_{M_{t'}}$ consisting of all micro-states x that between t' and t'' evolve, under the dynamics of the system, into macro-regions corresponding to macro-states of higher entropy: $\Gamma_{M_{t'}}^{(+)} := \{x \in M_{t'} \mid \phi_{t''-t'}(x) \in \Gamma_+\}$, where $\Gamma_+ := \bigcup_{M_i \in M_+} \Gamma_{M_i}$ and $M_+ := \{M_i \mid S_B(M_i) \geq S_B(M_{t'}), i = 1, \dots, m\}$. The probability that $S_B(t'') \geq S_B(t')$ then is equal to $p_{t'}(\Gamma_{M_{t'}}^{(+)}) = \mu(\Gamma_{M_{t'}}^{(+)})/\mu(\Gamma_{M_{t'}})$.

For Boltzmann's law to be true the condition that $\mu(\Gamma_{M_{t'}}^{(+)})/\mu(\Gamma_{M_{t'}}) \geq 1 - \varepsilon$, where ε is a small positive real number, must hold for all macro-states. Whether or not this is the case in a particular system is a substantial question, which depends on the system's dynamics. However, even if it is, 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 macro-state of higher entropy in the future, it is also overwhelmingly likely to have evolved into the current macro-state *from* a past macro-state M'' which also has *higher* entropy.¹³

Albert (2000: 71–96) discusses this problem at length and suggests fixing it 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 world came into being in whatever particular low-entropy highly condensed big-bang sort of macrocondition it is that the normal inferential procedures of cosmology will eventually present to us.

(Albert, 2000: 96)

This postulate is not without problems (see Winsberg (2004b) and Earman (2006) for discussions), and those who accept it disagree about its status in the theory (see Callender's and Price's contributions to Hitchcock (2004), which present opposite views). I assume that these issues can be resolved in one way or another and presuppose PH in what follows. The problem with wrong retrodictions can then be solved, so Albert suggests, by conditionalizing on PH:

[...] the probability distribution that one ought to plug into the equations of motion in order to make inferences about the world is the one that's uniform, on the standard measure, over those regions of the phase space of the world which are compatible both with *whatever* it is that we may happen to know about the present physical condition of the universe (just as the *original* postulate [SP]) *and* with the hypothesis that the original macrocondition of the universe was the one associated with the *big bang*.

(Albert, 2000: 95–96)

¹³ A point to this effect was first made by Ehrenfest and Ehrenfest (1907; 1912: 32–34). However, their argument is based on an explicitly probabilistic model and so its relevance to deterministic dynamical systems is tenuous.

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 macro-state 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 micro-state lies in A is

$$p_t(A | R_t) = \frac{\mu(A \cap R_t)}{\mu(R_t)} \quad (6.7)$$

where $R_t := M_t \cap \phi_{t-t_0}(M_p)$.

Again, whether or not BL is true given this postulate is a substantive question having to do both with the construction of the macro-states as well as the dynamics of the system; I will come back to this issue below. Let us first turn to the question of how these probabilities can be interpreted.

6.4.2 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) 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 consists of the trajectory of the system’s micro-state in phase space, on which the system’s macro-states 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 (2004), 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 then measured by the probability that it assigns to the actual course of history, i.e. by how likely it regards the things to happen that actually do happen. By definition, systems that do not involve probabilistic laws

have perfect fit. As an illustration, consider a Humean mosaic consisting only of ten outcomes of a coin flip: HHTHTTHHTT. 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 BSM as introduced above – the package of classical mechanics, PH and PHSP – is a putative best system of the sort just described and that PHSP probabilities can therefore be regarded as Humean chances:

Recall that (in Albert's formulation) the fundamental postulates of statistical mechanics are fundamental dynamical laws (Newtonian in the case of classical statistical mechanics), a postulate that the initial condition was one of low entropy, and the postulate that the probability distribution at the origin of the universe is the microcanonical distribution conditional on the low entropy condition. The idea then is that this package is a putative Best System of our world. The contingent generalisations it entails are laws and the chance statements it entails give the chances. It is simple and enormously informative. In particular, it entails probabilistic versions of all the principles of thermodynamics. That it qualifies as a best system for a world like ours is very plausible. By being part of the Best System the probability distribution earns its status as a law and is thus able to confer lawfulness on those generalisations that it (together with the dynamical laws) entails.

(Loewer, 2001: 618; cf. 2004: 1124)

There is an obvious problem with this view, namely that it assigns non-trivial probabilities (i.e. ones that can have values other than 0 and 1) to events within a deterministic framework, which Lewis himself thought was impossible (Lewis, 1986: 118). Loewer claims that Lewis was wrong on this and suggests that introducing probabilities via initial conditions solves the problem of reconciling determinism and chance:

[...] 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; cf. 2004: 1124)

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 modulo a qualification to which I turn below.

As above, I take the system's state at $t > t_0$ to be the macro-state M_t . We now need to determine the probability of the event 'being in set $A \subseteq \Gamma_{M_t}$ 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 A at time t ', $p_t(A)$, is equal to the probability of 'being in set A_0 at time t_0 ' where A_0 is, by definition, the set that evolves into A under the dynamics of the system after time t has elapsed. Formally, $p_t(A) = \mu_0(A_0) = \mu_0(\phi_{t_0-t}(A))$, where μ_0 is the micro-canonical distribution over the past state, i.e. $\mu_0(\cdot) = \mu(\cdot \cap \Gamma_{M_p})/\mu(\Gamma_{M_p})$.

The second part is 'conditionalizing on the entire macroscopic history of the world [...] up to time t '. Loewer does not define what he means by 'macroscopic history', but it seems natural to take a macro-history to be a specification of the system's macro-state at *every* instant of time between t_0 and t . A possible macro-history, for instance, would be that the system is in macro-state 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 macro-state into another. What we are now expected to calculate is the probability of 'being in set A at time t ' *given* the system's macro-history. Let Q_t be the set of all micro-states 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* macro-state of the system at t . The sought-after conditional probability then is $p_t(A | Q_t) = p_t(A \& Q_t)/p_t(Q_t)$, 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 rule introducing *L*-chances for deterministic systems, which I call the *L*-chance statistical postulate (LSP):

Let M_t be the system's macro-state at time t . SP is valid for the past state M_p , which obtains at t_0 . For all times $t > t_0$ the probability that the system's micro-state lies in A is

$$p_t(A | Q_t) = \frac{\mu_0(\phi_{t_0-t}(A \cap Q_t))}{\mu_0(\phi_{t_0-t}(Q_t))} \quad (6.8)$$

where Q_t is the subset of M_t of micro-states compatible with the entire past history of the system and, again, $\mu_0(\cdot) = \mu(\cdot \cap \Gamma_{M_p})/\mu(\Gamma_{M_p})$.

The crucial thing to realize now is that due to the conservation of the Liouville measure the expression for the conditional probability in PHSP can be expressed as $p_t(A | R_t) = \mu(\phi_{t_0-t}(A \cap R_t))/\mu(\phi_{t_0-t}(R_t))$. Trivially, we can substitute μ_0 for μ in this expression which makes it formally equivalent to Equation (6.8). Hence

PHSP can be interpreted as attributing probabilities to events at $t > t_0$ solely on the basis of the micro-canonical measure over the initial conditions, which is what Loewer needs.

However, the equivalence of PHSP and LSP is only formal because there is an important difference between Q_t and R_t : R_t only involves a conditionalization on PH, while Q_t contains the *entire* past history. This difference will be important later on.

6.4.3 Problems with fit

Loewer claims that BSM as introduced above is the system that strikes the best balance between simplicity, strength and fit. Trivially, this implies that BSM 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 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 (6.8) 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 (6.8). Of course, 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 the 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 discretize time by dividing the real axis into small intervals of length δ , then calculate 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 $p_t(A | Q_t)$ depended in a way on δ that would assure that the limit exists. This is not the case. And, surprisingly, the problem does not lie with the limit. It turns out that for all $t > t_1$ (i.e. after the first change of macro-state), the $p_t(A | Q_t)$ 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 macro-state 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 micro-states in Γ_{M_5} compatible with the past macro history. Now, these points must be such that they 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 an instant later are the ones that lie exactly on the boundary between M_1 and M_5 . But the boundary of a $6n$ dimensional region is $6n - 1$ dimensional and has measure zero. Therefore Q_t has measure zero for all $t > t_1$, and accordingly $p_t(A | Q_t)$ does not exist for $t > t_1$, no matter what A . Needless to say, this renders the limit $\delta \rightarrow 0$ obsolete.

The source of these difficulties seems to be Loewer's requirement that we conditionalize on the entire macro-history of the system. This suggests that the problem could be solved simply by reverting back to Albert's algorithm, which involves only conditionalizing on PH. Unfortunately the hope that this would make the problems go away is in vain. First, Loewer seems to be right that we need to conditionalize on the entire macro-history when calculating the fit of a system. Fit is defined as the probability that the system attributes to the entire course of actual history and hence the probability of an event at time t must take the past into account. It can turn out to be the case that the evolution of a system is such that the probability of an event at t is independent of (large parts of the) past of the system, in which case we need not take the past into account. But if certain conditional probabilities do not exist given the system's past, we cannot simply 'solve' the problem by ignoring it.

Second, even if one could somehow convincingly argue that conditionalizing on the entire past is indeed the wrong thing to do, this would not be the end of the difficulties because another technical problem arises, this time having to do with the limit. As above, let us discretize time by dividing the real axis into small intervals of length δ and calculate the relevant probabilities at the instants $t_0, t_0 + \delta, t_0 + 2\delta$ etc. The relevant probabilities are of the form $p(t', t'') := p(M_j \text{ at } t'' | M_i \text{ at } t' \ \& \ M_p \text{ at } t_0)$, where t' and t'' are two consecutive instances on the discrete time axis (i.e. $t'' > t'$ and $t'' - t' = \delta$), and M_i and M_j are the system's macro-states at t' and t'' respectively. Calculating these probabilities using PHSP we find:

$$p(t', t'') = \frac{\mu[\phi_{t'-t''}(\Gamma_{M_j}) \cap \Gamma_{M_i} \cap \phi_{t'-t_0}(\Gamma_{M_p})]}{\mu[\Gamma_{M_i} \cap \phi_{t'-t_0}(\Gamma_{M_p})]} \quad (6.9)$$

These probabilities exist and we then obtain the probability of the actual course of history by plugging in the correct M_i 's and multiply all of them.

The next step is to take the limit $\delta \rightarrow 0$, and this is where things go wrong. We need to distinguish two cases. First, the system is in the same macro-state at t' and t'' (i.e. $M_i = M_j$). In this case $\lim_{\delta \rightarrow 0} p(t', t'') = 1$ because $\lim_{\delta \rightarrow 0} \phi_{t'-t''}(\Gamma_{M_i}) = \Gamma_{M_i}$. Second, the system is in two different macro-states at t' and t'' . In this case $\lim_{\delta \rightarrow 0} p(t', t'') = 0$ because $\lim_{\delta \rightarrow 0} \phi_{t'-t''}(\Gamma_{M_j}) = \Gamma_{M_j}$ and, by definition, $\Gamma_{M_j} \cap \Gamma_{M_i} = \emptyset$. Hence the product of all these probabilities always comes out zero in the limit, and this for any phase flow ϕ and for any macro-history that involves at least one change of macro-state (i.e. for any macro-history in which the system ever assumes a state other than the past state). Fit, calculated in this way, fails to put different systems into a non-trivial fit ordering and is therefore useless.

The moral to draw from this is that we should not take the limit and instead calculate fit with respect to a finite number of instants of time; that is, we should calculate fit with respect to a discrete macro-history (DMH) rather than a continuous one. By a DMH I mean a specification for a finite number of instants of time $t_0 =: \tau_0 \leq \tau_1 \leq \dots \leq \tau_{j-1} \leq \tau_j := t_f$ (which can but need not be equidistant as the ones we considered above) of the system's macro-state at these instants: M_p at τ_0 , M_{τ_1} at τ_1 , M_{τ_2} at τ_2 , \dots , $M_{\tau_{j-1}}$ at τ_{j-1} and M_{τ_j} at τ_j , where M_{τ_i} is the system's macro-state at time τ_i .¹⁴

Before discussing the consequences of this move further, let me point out that once we go for this option, the seeming advantage of conditionalizing only on PH rather than the entire past history evaporates. If we only consider the system's macro-states at discrete times, Q_t no longer needs to have measure zero. As a result, the probabilities introduced in Equation (6.8) are well defined. So, conditionalizing on the entire past is no problem as long as the relevant history is discrete.

How could we justify the use of a discrete rather than a continuous macro-history? We are forced into conditionalizing over a continuum of events by the conjunction of three assumptions: (1) the posit that time is continuous, (2) the assumption that the transition from one macro-state to another one takes place at a precise instant, (3) the posit that fit has to be determined with respect to the *entire* macro-history of the system. We have to give up at least one of these to justify the use of a discrete rather than a continuous macro-history. The problem is that all

¹⁴ I 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.

three elements either seem reasonable or are deeply entrenched in the theory and cannot be renounced without far-reaching consequences.

The first option, discretizing time, would solve the problem because if we assume that time is discrete the macro-history is discrete too. The problem with this suggestion is that it is *ad hoc* (because time in classical mechanics is continuous), and therefore 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 macro-states; that is, allowing for there to be periods of time during which it is indeterminate in which macro-state the system is. This suggestion is not without merit as one could argue that sharp transitions between macro-states are indeed a mathematical idealization that is ultimately unjustifiable from a physics perspective. However, sharp transitions are a direct consequence of the postulate that macro-states supervene on micro-states. This postulate is central to the Boltzmannian approach and it is not clear how it could be given up without unsettling the entire edifice of BSM.

The third option denies that we should conditionalize 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 macro-state only at discrete instants and is oblivious about what happens between these. 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.¹⁵ From the point of view of Lewis' theory this seems unmotivated. Fit, like truth, is a semantic concept characterizing 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 is willing to believe that a discrete version of fit is satisfactory, it is not clear whether this leads to useful results. Depending on which particular instants of time one chooses to measure fit, one could get widely different results. These would be useful only if it was the case that the fit rankings came out

¹⁵ 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.

the same no matter the choice of instants. There is at least a question whether this is the case.

6.4.4 The putative best system is not the best system

I now assume, for the sake of argument, that the use of a DMH to calculate fit can be defended in one way or another. 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. As above, I first use Loewer's scheme and then discuss how conditionalizing *à la* Albert would change matters. The probability of the DMH is $p(DMH) = p_{\tau_0}(A_{\tau_0} | Q_{\tau_0}) \dots p_{\tau_{j-1}}(A_{\tau_{j-1}} | Q_{\tau_{j-1}})$, where the A_{τ_i} are those subsets of M_{τ_i} that evolve into $M_{\tau_{i+1}}$ under the evolution of the system between τ_i and τ_{i+1} . One can then prove that

$$p(DMH) = \mu_0[\Gamma_p \cap \phi_{\tau_0-\tau_1}(\Gamma_1) \cap \dots \cap \phi_{\tau_0-\tau_j}(\Gamma_j)] \quad (6.10)$$

where, for ease of notation, we set $\Gamma_p := \Gamma_{M_p}$ and $\Gamma_i := \Gamma_{M_{\tau_i}}$ for $i = 1, \dots, j$. Now define $N := \Gamma_p \cap \phi_{\tau_0-\tau_1}(\Gamma_1) \dots \phi_{\tau_0-\tau_j}(\Gamma_j)$. The fit of system is measured by the probability that it assigns to the actual DMH, which is given by Equation (6.10). 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 $\Gamma_p \setminus N$: $\mu_N(A) := k\mu_0(A \cap N)$, for all sets $A \subseteq \Gamma_p$ and $k = 1/\mu_0(N)$ (provided that $\mu_0(N) \neq 0$). Trivially, N contains the *actual* initial condition of the universe. An even simpler and more convenient distribution 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 maximize) 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 the demon, as, by assumption, he knows the entire Humean mosaic. One 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 μ_N or a Dirac delta function rather than μ_0 does not render the system much 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.

This conclusion can be opposed on various grounds. The first is that the introduction of alternative distributions makes the system radically more complex and hence on balance the original system is better. Consider the delta function first, and assume for the sake of argument that Γ_{M_p} is an interval. A counter to my claim then would be that the system with the delta function is complex because it has to store a number – the true initial condition – which is computationally costly because most numbers have decimal expansions that are infinitely long. By contrast, intervals can be stored cheaply; all the system has to store is something like ‘the unit interval’.

This argument is unconvincing because it relies on a misleading choice of examples. A specification of an interval demands the specification of two numbers, the interval’s boundaries. These *can*, of course, be natural numbers (and hence cheap to store), but this need not be the case; the boundaries of an interval can be real numbers. Moreover, it can also be the case that the true initial condition is a natural number, and hence cheap to store. So the bottom line is that storing an interval requires storing two numbers, while storing an initial condition only requires storing one number, and hence the claim that the system with the interval is simpler is unwarranted. Finally, this still leaves the question what happens if Γ_{M_p} is not an interval. Things are less clear in that case, but arguably storing a complicated boundary is more costly than storing a single number.

What about a distribution peaked over N ? The argument against this is that specifying N is more complicated than specifying Γ_{M_p} because the latter is given to us directly by the system’s macro-state structure, while the former has to be determined on the basis of the entire Humean mosaic in an extra step. While this is true, it does not seem to be correct to say that this is an ‘extra’ in the sense that it involves carrying out a step that we would not otherwise have to carry out. As becomes clear from Equation (6.10), in determining the fit of the system we effectively determine N , and there is no way around determining the system’s fit. So there is no extra work involved in determining N .

Another line of criticism claims that the gain in fit is not, as the above argument assumes, considerable.¹⁶ In fact, so the argument goes, the original system has already extremely good fit and fiddling around with the measure would result at best in a small improvement of fit; and this small improvement is not enough to compensate the loss in simplicity which is involved in using distributions other than the original one. Therefore, the system containing μ_0 is the best system after all.

I have already argued that I do not think that the use of alternative distributions makes the system more complex, but this criticism draws our attention to an

¹⁶ Thanks to Eric Winsberg for a helpful discussion on this issue.

important point well worth discussing, namely the absence of a specification in the putative best system of the Hamiltonian of the universe.¹⁷ Assume that the macro-history of the universe is indeed TD-like in that the entropy increases most of the time. Why then should we assume that the system consisting of Hamiltonian mechanics, PH and either PHSP or LSP has good fit? Whether or not this is the case depends on the dynamical properties of the Hamiltonian of the system. The phase flow ϕ – the solution of the equations of motion – explicitly occurs both in Equations (6.7) and (6.8) and hence whether these probabilities come out such that the system has good fit with the actual history depends on the properties of ϕ . In fact, we cannot calculate any probabilities at all if we do not have ϕ ! Without a phase flow both PHSP and LSP remain silent and it is absolutely essential that we are told what ϕ to plug into the probabilistic algorithms if we want to obtain any results at all. So it is somewhat puzzling that Loewer does not say anything about ϕ at all and Albert (2000: 96) completely passes over it in his definitive statement of his view (he says something about it at another place; I come back to this below).

We are now faced with the situation that we have to plug some ϕ into either PHSP or LSP in order to make predictions, and at the same time we are not told anything about which ϕ to take. One way out of this predicament would be to assume that any ϕ would do. That is, the claim would be that every classical system with PH and either PHSP or LSP behaves in such a way that the actual history comes out having high fit. This, however, is wrong. Whether a system behaves in this way, and whether the system has good fit with the actual history of the world, essentially depends both on the dynamics of the system and the structure of the macro-regions. Some Hamiltonians do not give rise to an approach to equilibrium at all (e.g. a system of harmonic oscillators), and others may have constants of motion whose invariant surfaces prevent the system from moving into certain regions of phase space into which it would have to move for the macro-history to be what it is.

Hence, the best system cannot possibly consist only of the above three elements; it also has to contain a specification of the macro-state structure and a Hamiltonian (from which we obtain ϕ by solving the equations of motion). If we assume the macro-state structure to be given by the combinatorial argument, then the challenge is to present a Hamiltonian (or a class of Hamiltonians) for which the fit of the system is high. How would such a Hamiltonian look?

¹⁷ Albert and Loewer regard Newtonian rather than Hamiltonian mechanics as part of the best system. Newtonian mechanics has a wider scope than Hamiltonian mechanics in that it allows also for dissipative force functions. However, such force functions do not occur in SM, where forces are assumed to be conservative. But the motions to which such forces give rise are more elegantly described by Hamilton mechanics. Hence I keep talking about Hamiltonian mechanics in what follows, which is not a drawback because the problems I describe would only be aggravated by considering full-fledged Newtonian mechanics.

Two criteria emerge from Albert's discussion. The first and obvious condition is that ϕ must be such that the probabilities calculated using PHSP come out right. The second, somewhat more involved, requirement is that ϕ be such that 'abnormal' micro-states (i.e. ones that lead to un-thermodynamic behaviour) are scattered in tiny clusters all over the macro-regions (Albert, 2000: 67, 81–85). The motivation for this condition is subtle and need not occupy us here; a detailed discussion can be found in Winsberg (2004a). Two things about these conditions are remarkable. First, Albert's discussion of these conditions suggests that they are technicalities that at some point have to be stated for completeness' sake and then can be put aside; in other words, his discussion suggests (although he does not say this explicitly) that these are simple conditions that hold true in all classical systems and once this fact is noticed we do not have to worry about them any more. This might explain why neither of them is mentioned in the final statement of his view (Albert, 2000: 96). This is misleading. These conditions are extremely special and many phase flows do not satisfy them. So it is very much an open question whether the Hamiltonians of those systems – in particular the universe as a whole – to which SM reasoning is applied satisfy these requirements.

To this one might reply that there is a sort of a transcendental argument for the conclusion that the Hamiltonian of the universe indeed must belong to this class of Hamiltonians, because if it did not then the universe simply would not be what it is and thermodynamics would not be true in it. The problem, and this is the second point, with this suggestion is that even if this argument was sound, it would not solve the original problem, namely that we cannot calculate any probabilities at all if we do not have ϕ . The fact that the Hamiltonian of the universe satisfies all required criteria does not tell us what this Hamiltonian – or the phase flow ϕ associated with it – is, and as long as we do not have the Hamiltonian we simply cannot use either PHSP or LSP. Hence a system that does not contain the Hamiltonian of the universe is not really a system at all because the probabilistic algorithms require the phase flow ϕ of the universe to be given.

As it stands we do not know what ϕ is and hence the package of classical mechanics, PH and either PHSP or LSP fails to produce predictions at all. But assume now, again for the sake of argument, that we somehow manage to get hold of the Hamiltonian (and can also solve the equations of motion so that we obtain ϕ). We could then say that the best system consists of Hamiltonian mechanics including the Hamiltonian of the universe, PH, the macro-state structure given by the combinatorial argument and a probabilistic algorithm. Before we can make this suggestion definitive we need to settle the question of which probabilistic algorithm to include. PHSP or LSP? There is again a trade-off between simplicity and fit. On the one hand, PHSP is simpler than LSP because conditionalizing on just the past state is simpler than conditionalizing on the entire discrete macro-history up

to time t . On the other hand, LSP will generally yield better predictions than PHSP and hence make the system stronger. Whether on balance PHSP or LSP is the better choice depends, again, on ϕ . It might be that the dynamics is so friendly that R_t and Q_t coincide, in which case PHSP is preferable on the grounds of simplicity. If this is not the case, then LSP may well be the better choice. Depending on which one comes out as striking the better balance, we add either PHSP or LSP to the above list. This then is our putative system.

Before raising a more fundamental objection against the view that the above quintet is a best system in the relevant sense, let me voice some scepticism about this line of defence against the original charge that fit can always be improved by substituting μ_0 by a measure like μ_N . The past state is defined *solely* via the value of some relevant macro-parameters and the corresponding macro-region, Γ_{M_p} , consists of *all* micro-states that are compatible with these values. In particular no facts about the behaviour of the $x \in \Gamma_{M_p}$ under the time evolution of the system is made. Now it would be something of a miracle if the overwhelming majority of points (relative to the Lebesgue measure) in a set *thus defined* would be the initial conditions of trajectories that are compatible with the *actual* course of history. We would have to have a very good and very special reason to assume that this is the case, and so far we do not have such a reason. And I doubt that there is one. For the approach to equilibrium to take place the system needs to be chaotic in some way, and such system involve divergence of nearby trajectories. For this reason it is plausible to assume that different $x \in \Gamma_{M_p}$ evolve into very different parts of phase space and thereby move through very different macro-regions, and hence give rise to different macro-histories. And for this reason using μ_N rather than μ_0 would improve fit.

There is a more fundamental objection to the view that the quintet of Hamiltonian mechanics, PH, the combinatorial macro-state structure, a Hamiltonian satisfying some relevant condition and, say, LSP is a best system in Lewis' sense. The point is that a system that does not contain probabilities at all is much simpler than one that does, at least if probabilities are introduced either by PHSP or LSP. The laws are made by Lewis' demon and in the deterministic world of classical mechanics, the simplest system that the demon can come up with is one that consists of Hamilton's equations together with the Hamiltonian of the system, the macro-state structure, and the exact initial condition – everything else follows from this. In other words, simply adding the true initial condition as an axiom to mechanics and the macro-state structure is much simpler than adding probabilities via either PHSP or SPP – and doing this is no problem for the demon as he knows the exact initial condition.

Before arguing for this let me locate the moot point. It is important to be clear on the fact that the ability to solve equations is irrelevant to the argument, because

both PHSP and LSP presuppose a solution to the equations of motion, because the phase flow ϕ , which is the solution of the equations of motion, explicitly figures both in Equations (6.7) and (6.8).¹⁸ In this respect systems comprising PHSP or LSP are not simpler than a system containing the exact initial condition as an axiom. So the controversy must be over the computational costs of handling initial conditions. The argument, then, must go something like this: having the laws and relevant macro-information is much simpler than knowing an initial condition, because storing an initial condition involves storing about $6n$ real numbers, which is costly because one has to take it in (presumably feed it into a computer), store it, and later on process it, which takes time and storage space and hence comes at a high cost of simplicity.

This is an illusion, which stems from the ambiguous use of ‘macro-information’. Macro-information – knowledge about the values of macroscopic parameters such as pressure and volume – about the past state is useless when we want to apply, say, PHSP. What goes into this probabilistic algorithm is not macro-information *per se*; it is the set of *micro-states* compatible with the available macro-information. Once this is realized, the apparent advantage of a system containing probabilities evaporates because taking in and storing information about an entire set is more costly than taking in and storing a single point. This is because specifying a set in $6n$ -dimensional space at least implies specifying its boundary, and this amounts to specifying a $(6n - 1)$ -dimensional hypersurface, containing infinitely many points. And when applying PHSP we have to evolve forward in time this entire surface, which involves evolving forward in time infinitely many points. How could that possibly be simpler than taking in, storing, and evolving forward one single point?

In sum, there are strong *prima facie* reasons to assume that the system’s fit can be improved by changing the initial measure. And even should it turn out that this is not the case, a system without probabilities would be better than one with probabilities. Therefore, neither PHSP nor LSP probabilities can be interpreted as Humean chances.

6.4.5 Interpreting BSM probabilities

Where does this leave us? I think that the above considerations make it plausible that what ultimately gives rise to the introduction of probabilities into classical mechanics are the epistemic limitations of those who use the theory. All we can know about a system is its macro-state and hence we put a probability distribution

¹⁸ Notice that this is a very strong, and entirely unrealistic, assumption. Characteristically we do not have solutions to the equations of motion of systems studied in SM, and many physicists see this as the motivation to use statistical considerations to begin with (see for instance Khinchin, 1949).

over the micro-states compatible with that macro-state, which then reflects our lack of knowledge.

How these epistemic probabilities should be understood is a question that needs to be further discussed. Here I can only hint at two possibilities. The first, a version of objective Bayesianism, appeals to Jaynes' 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 maximizes 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 systems. Hoefer's (2007) theory of Humean chance seems to make room for this possibility. This option should not scandalize the Humean. The bedrock of contemporary Humeanism as regards probability is the rejection of a metaphysical grounding of probabilities in propensities, tendencies, dispositions, and the like. At the same time the Humean regards probabilities as linked to human beliefs by the so-called Principal Principle, roughly the proposition that our subjective probabilities for an event to happen should match what we take to be its objective chance. Hence, for the Humean probabilities are both non-metaphysical and closely linked to beliefs. Hence, per se, the fact that SM probabilities have an epistemic component is no reason for outrage.

There are two main complaints about an epistemic interpretation of SM probabilities. The first¹⁹ 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. This argument has indeed some force when put forward against the Gibbs entropy if the SM probabilities are given an ignorance interpretation, because the Gibbs entropy is defined in terms of the probability distribution of the system. However, it has no force against an epistemic interpretation of PHSP or LSP probabilities simply because these probabilities do not occur in the Boltzmann entropy, which is defined in terms of the measure of certain parts of phase space (see Equation (6.2)). Probabilities simply have nothing to do with it.

Another frequently heard objection is a complaint about the alleged causal efficacy of human knowledge. The point becomes clear in the following – rhetorical – questions by Albert:²⁰

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

¹⁹ This point is often made in conversation, but I have been unable to locate it in print.

²⁰ Redhead (1995: 27–28, 32–33), Loewer (2001: 611), Goldstein (2001: 48), and Meacham (2005: 287–8) make essentially the same point.

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)²¹

It cannot be, and no one should think that it could. Proponents of epistemic probabilities need not believe in parapsychology, and therewith regard knowledge of tea and coffee as causally relevant to their cooling down.

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. There is no causal connection between knowledge and happenings in the world, and, at least in the context of classical SM, nothing of that sort is suggested by an epistemic interpretation of probabilities.

We have now reached the same point as in the discussions of macro-probabilities in the previous section. All that is needed to explain why things happen is the initial condition and the dynamics. *Prima facie* appearances notwithstanding, neither PHSP nor LSP probabilities have any role to play in explaining why a system behaves as it does. If these probabilities come out as BL would have them, this is indicative, it is a symptom, of the system behaving 'thermodynamically', but it is not the cause for this behaviour.

6.5 Conclusion

I have discussed two different ways of introducing probabilities into BSM and argued that the first one is irredeemably flawed, while the second leads to probabilities that are best understood as having an epistemic component. The discussions of both approaches have shown that there is a 'blind spot' in the literature on BSM, namely dynamics. Too much emphasis has been placed on probabilities and not enough attention has been paid to the dynamical conditions that have to fall in place for a system to behave TD-like. This puts two items on the agenda of future discussions of BSM: we need to understand better the nature of the epistemic

²¹ This contrasts starkly with the epistemic language that Albert uses throughout his discussion of PHSP. In the passages quoted above, for instance, he defines the past state as the one with which the 'normal inferential procedures of cosmology will eventually present us' and when discussing conditionalizing on PH he names 'make inferences about the world' as the aim and invites us to put a measure over those regions that are compatible with '*whatever* it is that we happen to know about the present physical condition of the universe'.

probabilities used in SM and we need to study more carefully the role that the dynamics of the system plays in explaining TD-like behaviour.²²

References

- Abraham, R. and Marsden, J. E. (1980). *Foundations of Mechanics*, 2nd edn. London: Benjamin-Cummings.
- Albert, D. (2000). *Time and Chance*. Cambridge, MA: Harvard University Press.
- Arnold, V. and Avez, A. (1968). *Ergodic Problems in Classical Mechanics*. New York, Benjamin.
- Boltzmann, L. (1877). Über die Beziehung zwischen dem zweiten Hauptsatze der mechanischen Wärmetheorie und der Wahrscheinlichkeitsrechnung resp. den Sätzen über das Wärmegleichgewicht. *Wiener Berichte*, **76**, 373–435. Reprinted in (1909). *Wissenschaftliche Abhandlungen*, vol. 2, ed. F. Hasenöhl. Leipzig: J. A. Barth, pp. 164–223.
- Callender, C. (1999) Reducing thermodynamics to statistical mechanics: the case of entropy. *Journal of Philosophy*, **96**, 348–373.
- Clark, P. (2001). Statistical mechanics and the propensity interpretation of probability. In *Chance in Physics: Foundations and Perspectives*. ed. J. Bricmont *et al.* Berlin: Springer-Verlag, pp. 271–281.
- Earman, J. (1986). *A Primer on Determinism*. Dordrecht: Kluwer.
- Earman, J. (2006). The ‘past hypothesis’: not even false. *Studies in History and Philosophy of Modern Physics*, **37**, 399–430.
- Earman, J. and Rédei, M. (1996). Why ergodic theory does not explain the success of equilibrium statistical mechanics. *British Journal for the Philosophy of Science*, **47**, 63–78.
- Ehrenfest, P. and Ehrenfest, T. (1907). Über Zwei Bekannte Einwände gegen das Boltzmannsche H-Theorem. *Physikalische Zeitschrift*, **8**, 311–14.
- Ehrenfest, P. and Ehrenfest, T. (1912/1959). *The Conceptual Foundations of the Statistical Approach in Mechanics*. Mineola, NY: Dover (Reprinted 2002. First published in German in 1912; first English Translation 1959.)
- Elga, A. (2004). Infinitesimal chances and the laws of nature. *Australasian Journal of Philosophy*, **82**, 67–76.
- Frigg, R. and Hoefer, C. (2007). Probability in GRW theory. *Studies in the History and Philosophy of Modern Physics*, **38**, 371–389.
- Frigg, R. (2008). A field guide to recent work on the foundations of statistical mechanics. In: *The Ashgate Companion to Contemporary Philosophy*. ed. D. Rickles. London: Ashgate, pp. 99–196.
- Goldstein, H. (1981). *Classical Mechanics*. Reading, MA: Addison Wesley.
- Goldstein, S. (2001). Boltzmann’s approach to statistical mechanics. In *Chance in Physics: Foundations and Perspectives*, ed. J. Bricmont *et al.* Berlin: Springer-Verlag.
- Goldstein, S. and Lebowitz, J. L. (2004). On the (Boltzmann) entropy of non-equilibrium systems. *Physica D: Nonlinear Phenomena*, **193** (1–4), 53–66.

²² Thanks to Jossi Berkovitz, Nancy Cartwright, David Lavis, Carl Hoefer, Barry Loewer, Eric Winsberg, 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.

- Hitchcock, C. (ed.) (2004). *Contemporary Debates in Philosophy of Science*. Oxford: Blackwell.
- Hoeyer, C. (2007). The third way on objective probability: a skeptic's guide to objective chance. *Mind*, **116**, 549–596.
- Khinchin, A. I. (1949). *Mathematical Foundations of Statistical Mechanics*. Mineola, NY: Dover.
- Lavis, D. (2005). Boltzmann and Gibbs: an attempted reconciliation. *Studies in History and Philosophy of Modern Physics*, **36**, 245–273.
- Lebowitz, J. L. (1993a). Boltzmann's entropy and time's arrow. *Physics Today*, September issue, 32–38.
- Lebowitz, J. L. (1993b). Macroscopic laws, microscopic dynamics, time's arrow and Boltzmann's entropy. *Physica A*, **194**, 1–27.
- Lebowitz, J. L. (1999). Statistical mechanics: a selective review of two central issues. *Reviews of Modern Physics*, **71**, 346–357.
- Lewis, D. (1986). A Subjectivist's Guide to Objective Chance and Postscripts to 'A subjectivist's guide to objective Chance'. In *Philosophical Papers*, vol. 2, Oxford: Oxford University Press, pp. 83–132.
- Lewis, D. (1994). Humean supervenience debugged. *Mind*, **103**, 473–90.
- Loewer, B. (2001). Determinism and chance. *Studies in History and Philosophy of Modern Physics*, **32**, 609–629.
- Loewer, B. (2004). David Lewis' Humean theory of objective chance. *Philosophy of Science*, **71**, 1115–1125.
- Meacham, C. (2005). Three proposals regarding a theory of chance. *Philosophical Perspectives*, **19**, 281–307.
- Redhead, M. (1995). *From Physics to Metaphysics*. Cambridge: Cambridge University Press.
- Sklar, L. (1993). *Physics and Chance. Philosophical Issues in the Foundations of Statistical Mechanics*. Cambridge: Cambridge University Press.
- van Lith, J. (2001). Ergodic theory, interpretations of probability and the foundations of statistical mechanics. *Studies in History and Philosophy of Modern Physics*, **32**, 581–594.
- Uffink, J. (2007). Compendium of the foundations of classical statistical physics. In *Philosophy of Physics*, ed. J. Butterfield and J. Earman. Amsterdam: North-Holland, pp. 923–1047.
- von Plato, J. (1988). Ergodic theory and the foundations of probability. In *Causation, Chance and Credence*, vol. 1, ed. B. Skyrms and W. L. Harper. Dordrecht: Kluwer, pp. 257–277.
- von Plato, J. (1989). Probability in dynamical systems. In *Logic, Methodology and Philosophy of Science*, vol. VIII, ed. J. E. Fenstad, I. T. Frolov and R. Hilpinen. Amsterdam: North-Holland, pp. 427–443.
- von Plato, J. (1994). *Creating Modern Probability*. Cambridge: Cambridge University Press.
- 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 statistical mechanics. *Philosophy of Science*, **71**, 707–718.