

Equilibrium in Boltzmannian Statistical Mechanics

Roman Frigg and Charlotte Werndl

Forthcoming in:
Eleanor Knox and Alastair Wilson (eds):
Routledge Companion to Philosophy of Physics

Contents

1	Introduction	3
2	Definitions of Equilibrium	5
3	The Approach to Equilibrium	9
4	Conclusion	14

1 Introduction

A gas that is confined to the left half of a container starts spreading as soon as the confining wall is removed and eventually spreads evenly over the entire available space. The state of being spread evenly is the *equilibrium state* and the process of expansion culminating in that state is the *approach to equilibrium*. It is one of the aims of statistical mechanics (SM) to give an exact characterisation of equilibrium, and to explain why and how systems approach the state of equilibrium, in terms of the dynamical laws that govern the individual molecules of which the gas is made up. What is it about molecules and their motions that leads them to spread out when the wall is removed? An important answer to these questions was suggested by Boltzmann (1909 [1877]), and variants of this answer are currently regarded by many as a promising option. In this chapter we introduce different versions of Boltzmannian SM (BSM) and discuss how they conceptualise equilibrium.¹

We begin by introducing the basic notions of BSM. A system in statistical mechanics has the mathematical structure of a *measure-preserving deterministic dynamical system* (X, μ, T_t) .² X is the state space and it contains all possible *micro-states* of the system. In the case of an isolated gas in a box made up of n particles a micro-state is given by the positions and the momenta of every molecule in the gas. The measure μ specifies how large certain parts of X are. We assume that the measure is normalised: $\mu(X) = 1$. This is for reasons of mathematical convenience and no connection between this measure and probability is assumed at this point. The time evolution of the system is given by the *evolution function* $T_t : X \rightarrow X$ where t is time and the function satisfies the requirement $T_{t_1+t_2}(x) = T_{t_2}(T_{t_1}(x))$ for all micro-states $x \in X$ and all instants of time t_1 and t_2 .³ Intuitively this means that it does not matter whether the time evolution is carried out in one go or takes place in stages: if the process begins in micro-state x , then we end up in the same final state $y \in X$ irrespective of whether we evolve x forward in time by $t_1 + t_2$ at once or whether we first evolve it x forward by t_1 and then evolve the resulting state forward by t_2 . The measure μ is assumed to be invariant under the dynamics, meaning that $\mu(T_t(A)) = \mu(A)$ for all measurable subsets A of X and all t .

¹Another approach that could be labelled ‘Boltzmannian’ departs from the so-called *Boltzmann equation*. For a discussion of this approach see Ardourel (2017), Brown, Myrvold and Uffink (2009), Uffink and Valente (2015) and Valente (2014).

²In line with most of the literature on BSM we focus on deterministic systems. For a discussion of BSM from a stochastic point of view see Werndl and Frigg (2017a).

³Time can either be continuous ($t \in \mathbb{R}$) or discrete ($t \in \mathbb{Z}$) (discrete time). Furthermore, it is a technical requirement that T_t is a measurable function in (t, x) . For a simple and intuitive introduction to dynamical systems see Berkovitz *et al.* (2016). Detailed discussions of the framework of BSM can be found in Frigg (2008) and Uffink (2007).

The macro-condition of a system is specified by its *macro-state*. In the case of the gas this can be done, for instance, by specifying the volume, pressure and temperature of the gas. It is usually assumed that there are a finite number m of macrostates, and we denote these by M_i , $i = 1, \dots, m$. The core posit of Boltzmannian statistical mechanics is that macro-states supervene on micro-states, meaning that a change in the macro-state must be accompanied by a change in the micro-state (it is not possible, for instance, to change the pressure of a gas without also changing the state of motion of at least some of its molecules). Hence, to every given micro-state x there corresponds *exactly one* macro-state. This determination relation is not one-to-one; usually many different micro-states correspond to the same macro-state. One can now group together all micro-states that belong to the same macro-state M_i . The result of this grouping is the *macro-region* X_{M_i} . All X_{M_i} taken together form a partition of X , meaning that they do not overlap and jointly cover X .

The *Boltzmann entropy* of a macro-state M_i is defined as $S_B(M_i) := k \log[\mu(X_{M_i})]$, where k is the Boltzmann constant; the Boltzmann entropy of a *system* at time t , $S_B(t)$, is the entropy of the macro-state of the system at t : $S_B(t) := S_B(M_{x(t)})$, where $x(t)$ is the micro-state at t and $M_{x(t)}$ is the macro-state supervening on $x(t)$.

Among the macro-states of the system two are of particular importance, namely the equilibrium state and the macrostate at the beginning of the process, also referred to as the ‘past state’. We introduce the special labels M_{eq} and M_p for these states (and choose a labelling of macro-states so that $M_p = M_1$ and $M_{eq} = M_m$). A crucial aspect of the standard presentation of BSM is that that the equilibrium macro-region $X_{M_{eq}}$ is vastly larger than any other macro-region. In fact $X_{M_{eq}}$ is so large that it takes up most of X .⁴ The size of the macro-region is generally seen as the crucial feature of equilibrium.

This raises two questions. First, why is equilibrium associated with the state that has the largest macro-region? The connection between equilibrium and large macro-regions is certainly not analytical: there is nothing in the *concept* of equilibrium tying it to the macro-state with the largest macro-region. Second, isolated systems, when left to themselves, end up in the equilibrium state. In our initial example the gas spreads until it fills the container evenly. Why do systems approach equilibrium? We discuss these problems in the next two sections.

⁴See, for instance, Albert (2000: 56-57), Bricmont (1995: 146); Goldstein (2001: 43, 45); Goldstein and Lebowitz (2004: 57), Penrose (1989: 403) and Zanghì (2005: 191, 196). However, Lavis (2005: 255-58; 2008: 685-87) points out that in some situations the equilibrium macro-region is *larger than any other macro-state* without taking up most of X . We set this problem aside, but it is discussed in Werndl and Frigg (2015b).

2 Definitions of Equilibrium

Why is the equilibrium macro-state the state with the largest macro-region? A prominent answer, now known as the *combinatorial argument*, originates in Boltzmann (1909 [1877]).⁵ The basic principle at work in Boltzmann’s argument is that equilibrium is the macro-state that is compatible with the largest number of micro-states. Based on this notion Boltzmann constructs X_{Meq} as follows. Consider the same system of n identical particles as above, but now focus on the 6-dimensional phase space χ of *one* of these particles (the space has six dimensions – the three position and three momentum coordinates of the particle). Only a finite portion χ_a of χ is accessible to the particle because the motion of the particle is constrained by the container walls and the total energy E of the gas. Now put a regular grid on χ_a such that the grid lines run parallel to the position and momentum axis, thereby dividing χ_a into a finite number of cells of equal size $\delta\omega$. This grid is also known as a coarse-graining of χ_a . Now label the grid cells ω_i for $i = 1, \dots, l$. The so-called coarse-grained micro-state of a particle is given by specifying in which cell ω_j its micro-state lies. The micro-state of the entire gas is a specification of the micro-state of every particle in the system, which is therefore determined by n labelled points in χ_a . The coarse-grained micro-state of the gas, also known as an *arrangement*, is a specification of which state of the particle lies in which cell of the partition.

The crucial observation now is that a number of arrangements correspond to the same macro-state because the macro-properties of the system are determined solely by the number of particles in each cell, while it is irrelevant exactly which particle is in which cell. For instance, whether particle number 5 and particle number 7 are in cells ω_1 and ω_2 respectively, or vice versa, makes no difference to the macro-properties of the system as a whole because these do not depend on which particle is in which cell. Hence, all one needs in order to determine the macro-properties of the system is a specification of how many particles there are in each cell of the coarse-graining of χ_a . Such a specification is called a *distribution* and it can be written as a tuple $D = (n_1, \dots, n_l)$, meaning that there are n_1 particles in cell ω_1 , etc. The n_j are referred to as *occupation numbers*. This allows us to ask a crucial question: how many arrangements are compatible with a given distribution D ? Some elementary combinatorial considerations show that

$$G(D) := \frac{n!}{n_1! \dots n_l!} \tag{1}$$

arrangements are compatible with a given distribution D (where ‘!’ denotes factorials, i.e. $k! := k(k-1) \dots 1$, for any natural number k and $0! := 1$). By construction

⁵Classical presentations of the argument can be found in Ehrenfest & Ehrenfest (2002 [1912]) and Tolman (1979 [1938]: Ch. 4). For discussions see Frigg (2008) and Uffink (2007).

$G(D)$ is a measure for how many micro states are compatible with a macro state: the larger the value of G , the more arrangements are compatible with a given distribution D .

The number $G(D)$ matters to our original question about the size of X_{Meq} . Since X is just a cartesian product of $6n$ copies of χ , each point in X corresponds to exactly one distribution D . Let us denote the region in X that corresponds to distribution D as X_D . One can then show that the size of that region is given by

$$\mu(X_D) = G(D) (\delta\omega)^n. \quad (2)$$

Hence the size of that part of X that corresponds to D is directly proportional to $G(D)$: the larger $G(D)$, the larger X_D . So by maximising $G(D)$ — that is, by finding the distribution that is compatible with the largest number of (coarse-grained) micro-states — one at the same also finds the largest macro-region. Let D_{max} be that distribution. If one then adopts the above principle that equilibrium is the macro-state that is compatible with the largest number of micro-states, then, by definition, D_{max} is the equilibrium distribution and $X_{D_{max}}$ is the equilibrium macro-region: $X_{D_{max}} = X_{Meq}$. Since D_{max} has the largest G of all distributions, it also corresponds to the largest macro-region, which provides the sought-after justification of the notion that equilibrium is the state with the largest macro-region.

It remains to find D_{max} . To solve this problem, Boltzmann makes two crucial sets of assumptions. The first concerns the energy of the particles. Boltzmann assumes that the energy of a particle only depends on which cell ω_j it is in, but *not* on the states of the other particles; that is, he neglects the contribution to the energy of the system that stems from interactions between the particles. He then also assumes that the energy of a particle only depends on which cell it is in but not on where it lies within the cell. Under these assumptions, the total energy of the system is given by $\sum_{j=1}^l n_j E_j$, where E_j is the energy of a particle in cell j . The second assumption is that there are many particles in each individual cell: ($n_j \gg 1$ for all j). Under these assumptions one can then prove that $G(D)$ reaches its maximum for

$$n_j = \alpha \exp(-\beta E_j), \quad (3)$$

which is the (discrete) *Maxwell-Boltzmann distribution*, where α and β are constants depending on the nature of the system. So the largest macro-region in X corresponds to the Maxwell-Boltzmann and that is the equilibrium region.

This justificatory strategy faces both conceptual and technical challenges. The main conceptual problem is the absence of a conceptual connection with the thermodynamic notion of equilibrium. The following is a typical TD textbook definition of

equilibrium: ‘A thermodynamic system is in equilibrium when none of its thermodynamic properties are changing with time [...]’ (Reiss 1996: 3). The problem is that there is no *conceptual* connection between this notion of equilibrium and the idea that the equilibrium macro-state is the one that is compatible with the largest number of micro-states. This is a problem for anyone wishing to maintain at least some congruence between SM and thermodynamics (TD).

Even if one was willing to set aside conceptual issues, there remains the technical problem that the assumptions made for the optimisation are so strong that the domain of application of the theory is in effect limited to dilute gases. The combinatorial argument assumes that the energy of a particle depends only on the cell in which it is located. This assumption applies, strictly speaking, *only to systems with non-interacting particles, i.e. ideal gases* (Frigg 2008; Uffink 2007). Ideal gases are, perhaps, a good approximation for *dilute gases*, i.e. gases of low density, and so the argument may deliver the approximately correct results for such systems. However, the argument remains silent about systems with stronger inter-particle forces such as liquids and solids. This is a serious limitation, and no suggestions have been made so far as to how it could be overcome.

One might try to avoid at least the technical difficulties by reading the argument backward, as it were: postulate that the Maxwell-Boltzmann distribution is the equilibrium distribution (rather than deriving it from an optimisation procedure), and then appeal to combinatorial considerations to establish that the corresponding macro-region is large. This move is motivated by the fact that Maxwell’s (1860) original derivation of the Maxwell-Boltzmann distribution does not appeal to optimisation procedures. Unfortunately, this is a blind alley. The Maxwell-Boltzmann distribution is in fact the equilibrium distribution only for a limited class of systems, namely for systems consisting of particles with negligible inter-particle forces. In general, systems with non-negligible interactions will have equilibrium distributions that are different from the Maxwell-Boltzmann distribution (Gupta 2003). A closer look at Maxwell’s derivation shows why: non-interaction enters via the postulate that the probability distributions in different spatial directions can be factorised, which is true only if there is no interaction between particles (see Uffink 2007).

A different justification appeals to the time-evolution of a system and argues that if a part of the state space is overwhelmingly large, then a state sooner or later has to move into that part and stay there for a long time. This, so the argument continues, is the defining feature of equilibrium, and therefore the largest macro-state is the equilibrium state. This reply is closely tied to the typicality approach that we discuss in the next section, where we will see that the basic assumptions of that

approach are questionable.

As we have seen in the last section, the Boltzmann entropy is proportional to the logarithm of the measure of a macro-region and therefore the macro-state with the largest macro-region also has the highest Boltzmann entropy. We know from TD that, if left to itself, a system approaches equilibrium, and equilibrium is the maximum entropy state. Therefore the macro-state with the largest macro-region is the equilibrium state.⁶

This line of argument faces a number of difficulties. The first is that it attributes an entropy to systems out of equilibrium and compares entropy levels at different stages of a process. But thermodynamics does not attribute an entropy to systems out of equilibrium at all and therefore comparing the entropies of different macro-states (most of which will be non-equilibrium states) makes little sense from a thermodynamic point of view. But even if this problem could be resolved (through a suitable generalisation of thermodynamics, for instance) there would remain a question why the fact that the thermodynamic entropy reaches a maximum in equilibrium would imply that this also holds for the Boltzmann entropy. To justify this inference, the assumption would need to be made that the thermodynamic entropy reduces to the Boltzmann entropy. However, it is far from clear whether that is the case. A connection between the TD entropy and the Boltzmann entropy has been established only for ideal gases where the so-called Sackur-Tetrode formula can be derived from BSM, which shows that both entropies have the same functional dependence on thermodynamic state variables. Yet for systems with interactions no such results are known (cf. Frigg and Werndl 2011b). Furthermore, there are well-known differences between the TD and the Boltzmann entropy. For example, the TD entropy is extensive but the Boltzmann entropy is not (Ainsworth 2012), and an extensive concept cannot reduce to a non-extensive concept (at least not without further qualifications).

A different route is taken in the long-run-fraction-of-time-account recently proposed by Werndl and Frigg (2015a, 2015b). This approach *defines* equilibrium in terms of how long a system spends in a certain state (rather than in terms of the size of its macro-region) and then proves a theorem establishing that the state is large in the requisite sense. Let $LF_M(x)$ be the long-run fraction of time that a system starting in micro-state x spends in macro-region X_M . If, for instance, $LF_M(x) = 0.43$, then the system starting in initial condition x spends 43% of the time in macro-state M in the long run. Equilibrium, then, is the macro-state in which the system spend

⁶This strategy has been mentioned to us in conversation but is hard to track down in print. Albert's (2000) considerations concerning entropy seem to gesture in the direction of this third strategy.

most of its time for nearly all initial conditions. Formally, if a macro-state M satisfies the following condition, then it is the system's equilibrium state: $LF_M(x) \geq \alpha$ for a real number $\alpha \in (0.5, 1]$ and for all $x \in Y$ where Y is a subset of X such that $\mu(Y) \geq 1 - \varepsilon$ for a very small real number $\varepsilon \geq 0$. An obvious question concerns the value of α . Often the assumption seems to be that α is close to one. This is a reasonable but not the only possible choice, and nothing hangs on the value of α . The introduction of ε accounts for the possibility that we should not expect *every* initial condition to approach equilibrium (Callender 2001).

This definition is couched entirely in terms of time and remains completely silent about the size of X_M . This raises the question: is X_M large in the way BSM would have it? The answer to this question is given by the Dominance Theorem: If M is an equilibrium, then $\mu(X_M) \geq \alpha(1 - \varepsilon)$ (Werndl and Frigg 2015b). The theorem is completely general in that no dynamical assumptions are made (in particular it is not assumed that the system is ergodic) and hence the theorem also applies to strongly interacting systems such as solids and liquids. It is important to note that the theorem makes the conditional claim that *if* there is an equilibrium, *then* X_M is large. As with all conditionals, the crucial and often vexing question is whether, and under what conditions, the antecedent holds.

3 The Approach to Equilibrium

Why do isolated systems, when left to themselves, approach equilibrium? To appreciate the thrust of this question we have to recall an important property of the phenomenology of the approach to equilibrium: irreversibility. We see gases spread, yet we never observe the reverse process of a uniformly distributed gas suddenly concentrating in the left half of the container. So it seems that the question ought to be: why do systems approach equilibrium *irreversibly*? The standard line is that this irreversible approach to equilibrium is a consequence of the second law of thermodynamics, which, roughly, states that entropy cannot decrease in isolated systems. However, as Brown and Uffink (2001) rightly point out, that systems prepared in a non-equilibrium state do approach equilibrium is not a consequence of the second law (or any other law of standard thermodynamics) and has to be added as an independent principle, which they call the 'minus first law'. This suggests that the task for non-equilibrium SM is to derive the minus first and the second law of thermodynamics from BSM.

This is setting the bar too high. SM systems show Poincaré recurrence. That is, the system's micro-state will eventually return arbitrarily close to the original initial

condition.⁷ But a system with Poincaré recurrence cannot possibly exhibit strict irreversible behaviour. Contributors to the discussion have acknowledged this fact (more or less explicitly) and there is an (at least tacit) agreement that what should be derived from SM is an *approximate* version of the second law. Different versions of BSM differ in how they explicate the approximation.

Boltzmann (1909 [1877]) conceptualised the problem as one of showing that the system is overwhelmingly likely to be in equilibrium whenever one observes the system. To this end Boltzmann defined macrostates in terms of distributions (he assumed that to every distribution D there corresponds a macro-state) and introduced the postulate that the probability of a macro-state is proportional to $G(D)$. Since, as we have seen above, $G(D)$ is largest for the equilibrium state, it follows immediately that equilibrium is the most likely state and he argued that ‘the system of bodies always evolves from a more unlikely to a more likely state’ (Boltzmann 1909 [1877]: 166; our translation). This view faces a number of problems (Frigg 2010a). There is the question of how to justify the postulate that the probability of a macro-state is proportional to $G(D)$, and, more seriously, it is unclear where the tendency to move towards more probable states comes from. The probabilities of macro-states 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.

A related approach (also originating in the work of Boltzmann) appeals to the notion of ergodicity. Intuitively, a system is ergodic if the fraction of time an arbitrary solution stays in A equals the measure of A in the long run. More formally, let $LF_A(x)$ be the long run fraction of time a system starting in initial condition x will spend in a set $A \subset X$ (that is, for time $t \rightarrow \infty$). A system is ergodic iff $LF_A(x) = \mu(A)$ for all subsets $A \subset X$ and for almost all initial conditions x (that is, except perhaps for some initial conditions that taken together form a set of measure zero). If, for instance, a certain set has measure $1/3$ and if the system is ergodic, then we know that in the long run it will spend $1/3$ of the time in that set.

Since ergodicity concerns any set $A \subset X$, it a fortiori also concerns macro-regions, and hence an ergodic system will spend a fraction of time in every macro-state that is proportional to the size of the corresponding macro-region. Since equilibrium is the state with the largest macro-region, an ergodic system will spend most of the time in equilibrium for almost all initial conditions (which implies, of course, that if system is set in motion in a non-equilibrium micro-condition, it will soon enough move into

⁷That this may take a very long time to happen is besides the point as far as a justification of the second law is concerned.

the equilibrium macro-region). This is the justification of the ergodic programme of the minus first and the second laws of thermodynamics. It is an approximate justification in two ways. First, systems are said to spend most of the time in equilibrium but they are not said to never move out of equilibrium. In fact they can (and they will), and so ergodicity does not imply strict irreversibility. Second, ergodicity does not rule out that there are ‘bad’ initial conditions, i.e. initial conditions that lie on trajectories that do not satisfy $LF_A(x) = \mu(A)$; ergodicity just requires that there are ‘few’ bad initial conditions in the sense that all bad conditions taken together form a set of measure zero.

The ergodic programme faces two main challenges. The first, known as the *measure zero problem*, points out that the set of allowed ‘bad’ states can actually rather large since measure zero sets can be rather big. As an example consider the set of the rational numbers. This set has measure zero in the real numbers, but there are great many rational numbers. This fact can be made visible by using alternative ways to assess the size of set, for instance Baire categories (Sklar 1993: 182-88). The second objection, the *irrelevancy challenge*, is that ergodicity is irrelevant because many real systems are not ergodic (Earman and Rédei 1996). The force of this argument can be mitigated by introducing the notion of epsilon-ergodicity (Vranas 1998). Intuitively, a dynamical system is epsilon-ergodic iff it is ergodic on the vast majority of X , namely on a set of measure $\geq 1 - \varepsilon$, where ε is very small real number or zero. The class of systems that are epsilon-ergodic is larger than the class of systems that is ergodic and it is plausible that it comprises many realistic systems such as gases and some liquids (Frigg and Werndl 2011a). However, there will be SM systems that are not epsilon-ergodic and it remains unclear how the ergodic programme deals with them.

An alternative approach is developed by Albert (2000), who proposes to explain irreversible behaviour in terms of transition probabilities: rather than assigning probabilities to macro-states, we should look at how likely a system is to transition into a right macro-state given that it is in certain macro-state now. To this end Albert introduces the *statistical postulate*: given that the system is in macro-state M at time t , the probability of finding the micro-state of the system in a set $C \subset X_M$ at time t is $\mu(C)/\mu(X_M)$. One can separate the states in X_M into ‘good’ and ‘bad’ ones, with good ones being those lying on trajectories that move into macro-states of higher entropy (than M) once they leave X_M and bad ones being those lying on trajectories that move into lower entropy (than M) states. Thermodynamic behaviour is then justified if one can show that for all macro-states (other than the equilibrium macro-state) it is the case that the system is overwhelmingly more likely to move toward a higher (Boltzmann) entropy macro-state than to a lower (Boltzmann)

entropy macro-state. However, the statistical postulate also allows to calculate the probability that a system has moved into the current macro-state from a higher or a lower entropy macro-state. It turns out to be the case that whenever 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 which also has higher entropy. So the formalism returns wrong transition probabilities.

Albert 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*, 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). The problem with the wrong transition probabilities is then solved by conditioning on the past state: given that the system is in macro-state M at time t , the probability of finding the system’s micro-state in set $C \subset X_M$ at time t is $\mu(C \cap T_t(X_{M_p}))/\mu(X_M \cap T_t(X_{M_p}))$. Albert then argues at length that if this rule is used to calculate probabilities, then a high entropy future as well as a low entropy past are overwhelmingly likely.

There are number of concerns about this explanation. A crucial aspect of Albert’s explanation of the second law is to apply SM to the universe as a whole. Earman argued that this project is doomed to failure because in the setting of standard cosmological models the past hypothesis is ‘not even false’ (Earman 2006: 400).⁸ The Boltzmann entropy is a global quantity characterising the macro-state of an entire system, in Albert’s case the entire universe.

As Winsberg (2004) point out, just because the overall entropy of the universe increases, it need not be the case that the entropy in a small subsystem also increases and hence Albert’s approach cannot explain the behaviour of small systems like gases in laboratories. Albert’s argument relies on assigning probabilities to sets of microstates based on a certain algorithm, and there are questions about the justification of that algorithm (see Davey 2008 and Frigg 2010a). Furthermore, Albert’s justification that a high entropy future and low entropy past are overwhelmingly likely appeals to a dynamical assumption, which he calls the ‘scattering condition’, and there is a question whether this condition holds true in SM systems.⁹

⁸For a further discussion of the past hypothesis see Wallace (forthcoming), and for a discussion of its explanatory relevance see see Callender (2004) and Price (2004)

⁹For further discussions see chapter 7.3 on The Entropy Asymmetry and Frigg (2008, Sec. 3.2.5)

An alternative account explains the approach to equilibrium in terms of *typicality*. Consider an element e of a set Σ . Typicality is a relational property of e , which e possesses with respect to Σ , a property P and a measure ν , often referred to as the ‘typicality measure’. Intuitively, e is typical if most members of Σ have property P and e is one of them. More precisely, let $\Pi \subset \Sigma$ be the set of all elements that have property P . Then the element e is typical iff $e \in \Pi$ and $\nu(\Pi)/\nu(\Sigma) \geq 1 - \varepsilon$, where ε is a finite but very small positive real number. The element of interest in SM is a micro-state x and the typicality measure is the Lebesgue measure μ . The typicality account comes in different versions that disagree about the choice of the set Σ and the selection of property P (for a discussion of these different accounts see Frigg (2009) and (2010b)).

The first account explains the approach to equilibrium in terms of equilibrium micro-states being typical. As we have seen above, the equilibrium macro-region is by far the largest macro-region. This implies that equilibrium micro-states are typical with respect to X because ‘reaching the equilibrium distribution in the course of the temporal evolution of a system is inevitable due to the fact that the overwhelming majority of microstates in the phase space have this distribution’ (Zanghì 2005: 196; our translation). However, as Uffink (2007: 979-80) points out, in general there is no reason to assume that points in an atypical set have to evolve into a typical set; typical states are not attractors of atypical states. The second account combines a typicality claim about equilibrium micro-states with a further typicality claim about the dynamics of the system. Goldstein champions such an account when he submits that for ‘a non-equilibrium phase point $[x]$ of energy E , the Hamiltonian dynamics governing the motion $[x(t)]$ would have to be ridiculously special to avoid reasonably quickly carrying $[x(t)]$ into $[X_{M_{eq}}]$ and keeping it there for an extremely long time — unless, of course, $[x]$ itself were ridiculously special.’ (Goldstein 2001: 43-44). Unfortunately, Goldstein offers no account of what it means for a dynamical law not to be ‘ridiculously special’ and so the account remains underspecified (see Frigg and Werndl 2012 for proposed completion of the account in terms of epsilon-ergodicity.) The third account, due to Lebowitz (1993a, 1993b), considers the internal structure of macro-regions (in much the same way as the transition probability approach we have seen above) and argues that micro-states lying on entropy-increasing trajectories are typical. This line of argument faces the same challenge as Albert’s and it remains an unproven assertion that the internal structure of the macro-regions is as the account asserts (see Frigg 2010b for a discussion of this issue).

The problem of the approach to equilibrium takes a different form in the long-run-fraction-of-time-account (Werndl and Frigg 2015a, 2015b). In that account it is part and parcel of the notion of an equilibrium state that the system approaches that state and stays there most of the time; if a state does not have that feature, then it

is not an equilibrium state. The crucial question then is: under what conditions does an equilibrium exist? Werndl and Frigg (2017b) point out that for an equilibrium to exist three factors need to co operate: the choice of macro-variables, the dynamics of the system, and the choice of the effective state space. They then prove a theorem providing general necessary and sufficient conditions for the existence of an equilibrium state. Intuitively, the theorem says that there is an equilibrium just in case the effective state space of the system is split up into invariant regions on which the motion is ergodic and the equilibrium macro-state takes up most of each region. This suggests a change in the way in which we think about equilibrium: rather than launching a search for one crucial factor (such as ergodicity or typicality), the focus should be on finding triplets of macro-variables, dynamical conditions, and effective state spaces that satisfy the conditions of the theorem. An example of such triplet is the dynamics of the Kac ring on the full state space with a homogeneity macro-variable (Werndl and Frigg 2015b). With this theorem the theoretical question when and under what conditions an equilibrium exists is solved. But the identification and classification of these triplets for concrete problems remains an open question.

4 Conclusion

We have reviewed a number of Boltzmannian definitions of equilibrium along with explanations of the approach to equilibrium. We want to conclude by drawing attention to two further issues. The first is the relation of BSM to Gibbsian SM (see chapter 7.2 on Equilibrium in Gibbsian Statistical Mechanics). There are two competing theoretical approaches in SM, one associated with Boltzmann and the other with Gibbs, which offer different conceptualisations of equilibrium. This would not be a problem if the two formalisms were equivalent (for instance, in a similar way in which the Schrödinger and the Heisenberg picture in quantum mechanics are equivalent). Unfortunately they are not, and there is no obvious way to translate results from one framework into the other. But having two incompatible notions of equilibrium at work in SM is unsatisfactory and a sustained reflection on how the Boltzmannian and the Gibbsian approach relate to each other is necessary. Steps towards a better understanding of the relationship between these two approaches are made in Lavis (2005) and Werndl and Frigg (2017c), but the problem is one that deserves more attention that it has received so far.

The second issue is the interpretation of probability. As we have seen, explanations of the approach to equilibrium all rely in one way or another on probabilities, and there is question of how these should be interpreted. A number of suggestions have been made including Humean Best Systems (Frigg and Hoefer 2015, Loewer 2001), typicality measures (Werndl 2013), frequencies (discussed but not endorsed in van

Lith 2001) and propensities (discussed but not endorsed in Clark 2001). Another approach attempts to ground the probabilities of SM in quantum mechanics, and so they are interpreted in whatever way quantum probabilities are interpreted. See chapter 7.4 on Quantum Foundations for a discussion of these approaches. There is no consensus on this issue, and finding a coherent interpretation of SM probabilities remains a challenge.

References

- Ainsworth, P. M. (2012) “Entropy in Statistical Mechanics,” *Philosophy of Science* 79: 542-60.
- Albert, D. (2000) *Time and Chance*, Cambridge/MA and London: Harvard University Press.
- Ardourel, V. (2017) “Irreversibility in the derivation of the Boltzmann equation,” *Foundations of Physics* 47: 471-89.
- Berkovitz, J. *et al.* (2016) “The Ergodic Hierarchy,” *Stanford Encyclopedia of Philosophy* (Summer 2016 Edition), Edward N. Zalta (ed), URL = <https://plato.stanford.edu/archives/sum2016/entries/ergodic-hierarchy/>.
- Boltzmann, L. (1909 [1877]) “Über die Beziehung zwischen dem zweiten Hauptsatze der mechanischen Wärmetheorie und der Wahrscheinlichkeitsrechnung resp. den Sätzen über das Wärme Gleichgewicht,” in F. Hasenöhl (ed), *Wissenschaftliche Abhandlungen Leipzig* 2, pp. 164-223.
- Brown, H. and Uffink, J. (2001) “The origin of time-asymmetry in thermodynamics: the minus first law,” *Studies in History and Philosophy of Modern Physics* 32: 525-38.
- Brown, H., Myrvold, W. and Uffink, J. (2009) “Boltzmann’s H-theorem, its discontents, and the birth of statistical mechanics,” *Studies in History and Philosophy of Modern Physics* 40: 174-91.
- Bricmont, J. (1996) “Science of chaos or chaos in science?,” in P. R. Gross, N. Levitt and M. W. Lewis (eds), *The Flight from Science and Reason* (Annals of the New York Academy of Sciences, Vol. 775), New York: The New York Academy of Sciences, pp. 131-75.

- Callender, C. (2001) "Taking thermodynamics too seriously," *Studies in the History and Philosophy of Modern Physics* 32: 539-53.
- (2004) "There is no puzzle about the low-entropy past," in C. Hitchcock (ed), *Contemporary Debates in Philosophy of Science*, Oxford, Malden/MA and Victoria: Blackwell, pp. 240-55.
- Clark, P. (2001) "Statistical mechanics and the propensity interpretation of probability," in J. Bricmont, D. Durr, M. C. Galavotti, G. Ghirardi, F. Petruccione and N. Zanghì (eds), *Chance in Physics: Foundations and Perspectives* (Lecture Notes in Physics, Volume 574), Berlin, Heidelberg and New York: Springer, pp. 271-81.
- Davey, K. (2008) "The Justification of Probability Measures in Statistical Mechanics," *Philosophy of Science* 75: 28-44.
- 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. (2002 [1912]) *The Conceptual Foundations of the Statistical Approach in Mechanics*, Mineola/New York: Dover Publications.
- Frigg, R. (2008) "A Field Guide to Recent Work on the Foundations of Statistical Mechanics," in D. Rickles (ed), *The Ashgate Companion to Contemporary Philosophy of Physics*, London: Ashgate, pp. 99-196.
- Frigg, R. (2009) "Typicality and the Approach to Equilibrium in Boltzmannian Statistical Mechanics," *Philosophy of Science* 76: 997-1008.
- Frigg, R. (2010a) "Probability in Boltzmannian Statistical Mechanics," in G. Ernst and A. Hüttemann (eds), *Time, Chance and Reduction. Philosophical Aspects of Statistical Mechanics*, Cambridge: Cambridge University Press, pp. 92-118.
- Frigg, R. (2010b) "Why Typicality Does Not Explain the Approach to Equilibrium," in M. Suárez (ed), *Probabilities, Causes and Propensities in Physics*, Dordrecht: Springer, pp. 77-93.
- Frigg, R. and Hoefer, C. (2015) "The Best Humean System for Statistical Mechanics," *Erkenntnis* 80: 551-74.

- Frigg, R. and Werndl, C. (2011a) “Explaining Thermodynamic-Like Behaviour in Terms of Epsilon-Ergodicity,” *Philosophy of Science* 78: 628-52.
- Frigg, R. and Werndl, C. (2011b) “Entropy - A Guide for the Perplexed,” in C. Beisbart and S. Hartmann (eds), *Probability in Physics*, Oxford: Oxford University Press, pp. 115-42.
- Frigg, R. and Werndl, C. (2012) “Demystifying Typicality,” *Philosophy of Science* 79: 917-29.
- Goldstein, S. (2001) “Boltzmann’s approach to statistical mechanics,” in J. Bricmont, D. Durr, M. C. Galavotti, G. Ghirardi, F. Petruccione and N. Zanghi (eds), *Chance in Physics: Foundations and Perspectives* (Lecture Notes in Physics, Volume 574), Berlin, Heidelberg and New York: Springer, pp. 39-54.
- Goldstein, S. and Lebowitz, J. (2004) “On the (Boltzmann) entropy of nonequilibrium systems,” *Physica D* 193: 53-66.
- Gupta, M. C. (2003) *Statistical Thermodynamics*, New Delhi: New Age.
- Lavis, D. (2005) “Boltzmann and Gibbs: An attempted reconciliation,” *Studies in History and Philosophy of Modern Physics* 36: 245-73.
- (2008) “Boltzmann, Gibbs and the Concept of Equilibrium,” *Philosophy of Science* 75: 682-96.
- Loewer, B. (2001) “Determinism and chance,” *Studies in History and Philosophy of Modern Physics* 32: 609-29.
- Lebowitz, J. (1993a): “Boltzmann’s Entropy and Time’s Arrow,” *Physics Today* 46: 32-38.
- (1993b) “Macroscopic Laws, Microscopic Dynamics, Time’s Arrow and Boltzmann’s Entropy,” *Physica A* 194: 1-27.
- Maxwell J. C. (1860). “Illustrations of the dynamical theory of gases. Part 1. On the motions and collisions of perfectly elastic spheres; Part 2. On the process of diffusion of two or more kinds of moving particles among one another,” *Phil. Mag.*, 19, 20: 19-32, 21-37. ser. 4, 19.
- Penrose, R. (1989) *The Emperor’s New Mind*, Oxford: Oxford University Press.
- Price, H. (2004) “On the origins of the arrow of time: Why there is still a puzzle about the low-entropy past,” in C. Hitchcock (ed), *Contemporary Debates in Philosophy of Science*, Oxford and Malden/MA: Blackwell, pp. 219-39.

- Reiss, H. (1996) *Methods of Thermodynamics*, Mineola/NY: Dover.
- Sklar, L. (1993). *Physics and Chance. Philosophical Issues in the Foundations of Statistical Mechanics*, Cambridge: Cambridge University Press.
- Tolman, R. C. (1979 [1938]) *The Principles of Statistical Mechanics*, Mineola/New York: Dover.
- Uffink, J. (2007) “Compendium of the foundations of classical statistical physics,” in J. Butterfield and J. Earman (eds), *Philosophy of Physics*, Amsterdam: North Holland, pp. 923-1047.
- Uffink J. and Valente, G. (2015) “Lanford’s Theorem and the Emergence of Irreversibility,” *Foundations of Physics* 45: 404-38.
- Valente G. (2014) “The approach towards equilibrium in Lanford’s theorem,” *European Journal for Philosophy of Science* 4: 309-35.
- van Lith, J. (2001) “Ergodic theory, interpretations of probability and the foundations of statistical mechanics,” *Studies in History and Philosophy of Modern Physics* 37: 581-94.
- Vranas, P. B. M. (1998) “Epsilon-ergodicity and the success of equilibrium statistical mechanics,” *Philosophy of Science* 65: 688-708.
- Wallace, D. (forthcoming) “The Nature of the Past Hypothesis,” forthcoming in Proceedings of the conference Philosophy of Cosmology, Tenerife, 2014.
- Werndl, C. (2013) “Justifying Typicality Measures in Boltzmannian Statistical Mechanics,” *Studies in History and Philosophy of Modern Physics* 44: 470-79.
- Werndl C. and Frigg, R. (2015a) “Rethinking Boltzmannian Equilibrium,” *Philosophy of Science* 82: 1224-35.
- Werndl C. and Frigg, R. (2015b) “Reconceptualising Equilibrium in Boltzmannian Statistical Mechanics and Characterising its Existence,” *Studies in History and Philosophy of Modern Physics* 49: 19-31.
- Werndl C. and Frigg, R. (2017a) “Boltzmannian Equilibrium in Stochastic Systems,” forthcoming in M. Massimi and J.-W. Romeijn (eds), *Proceedings of the EPSA15 Conference*, Berlin and New York: Springer.
- Werndl C. and Frigg, R. (2017b) “When Does a Boltzmannian Equilibrium Exist?,” forthcoming in D. Bedingham, O. Maroney and C. Timpson (eds), *Quantum Foundations of Statistical Mechanics*, Oxford: Oxford University Press.

- Werndl C. and Frigg, R. (2017c) “Mind the Gap: Gibbs versus Boltzmann,” forthcoming in *Philosophy of Science*.
- Winsberg, E. (2004) “Can conditioning on the past hypothesis militate against the reversibility objections?,” *Philosophy of Science* 71: 489-504.
- Zanghì, N. (2005) “I fondamenti concettuali dell’approccio statistico in fisica,” in V. Allori, M. Dorato, F. Laudisa and N. Zanghì (eds), *La Natura Delle Cose. Introduzione ai Fundamenti e alla Filosofia della Fisica*, Roma: Carocci.