Rethinking Boltzmannian Equilibrium

Charlotte Werndl and Roman Frigg*

Boltzmannian statistical mechanics partitions the phase space of a system into macroregions, and the largest of these is identified with equilibrium. What justifies this identification? Common answers focus on Boltzmann's combinatorial argument, the Maxwell-Boltzmann distribution, and maximum entropy considerations. We argue that they fail and present a new answer. We characterize equilibrium as the macrostate in which a system spends most of its time and prove a new theorem establishing that equilibrium thus defined corresponds to the largest macroregion. Our derivation is completely general and does not rely on assumptions about the dynamics or interparticle interactions.

1. Introduction. Boltzmannian statistical mechanics (BSM) partitions the phase space of a system into cells consisting of macroscopically indistinguishable microstates. These cells correspond to the macrostates, and the largest cell is singled out as the equilibrium macrostate. The connection is not conceptual: there is nothing in the concept of equilibrium tying equilibrium to the largest cell. So what justifies the association of equilibrium with the largest cell?

After introducing BSM (sec. 2), we discuss three justificatory strategies based on Boltzmann's combinatorial argument, the Maxwell-Boltzmann distribution, and maximum entropy considerations, respectively. We argue that all three fail because they either suffer from internal difficulties or are restricted to systems with negligible interparticle forces (sec. 3). This prompts the search for an alternative answer. In analogy with the standard thermodynamic definition of equilibrium, we characterize equilibrium as the macrostate in which the system spends most of its time. We then present a new mathematical theorem proving that such an equilibrium macrostate indeed corresponds to the largest cell (sec. 4). This result is completely general in

*To contact the authors, please write to: Department of Philosophy, Logic and Scientific Method, London School of Economics, Houghton Street, WC2A 2AE, UK; Charlotte Werndl, e-mail: charlotte.werndl@sbg.ac.at; Roman Frigg, e-mail: r.p.frigg@lse.ac.uk.

Philosophy of Science, 82 (December 2015) pp. 1224–1235. 0031-8248/2015/8205-0041$10.00

that it is not based on any assumptions about the system's dynamics or the nature of interactions within the system.

2. Boltzmannian Statistical Mechanics. Let us briefly introduce BSM.[1] Consider a system consisting of n particles that is isolated from the environment and in a bounded container. The system's state is specified by a point $x = (q, p)$ (the *microstate*) in its $6n$-dimensional phase space Γ. The system's dynamics is determined by its classical Hamiltonian $H(x)$. Energy is preserved, and therefore the motion is confined to the $6n - 1$ dimensional energy hypersurface Γ_E defined by $H(x) = E$, where E is the energy value. The solutions of the equations of motion are given by the phase flow ϕ_t on Γ_E, where $\phi_t(x)$ is the state into which $x \in \Gamma_E$ evolves after t time steps. The Lebesgue σ-algebra is Σ_E and, intuitively speaking, consists of all relevant subsets of Γ_E. Space Γ is endowed with the Lebesgue measure μ, which is preserved under ϕ_t. This measure can be restricted to a measure μ_E on Γ_E, which is preserved as well and is normalized; that is, $\mu_E(\Gamma_E) = 1$. The quadruple $(\Gamma_E, \Sigma_E, \mu_E, \phi_t)$ is a *measure-preserving dynamical system*.

Assume that the system can be characterized by a set $\{v_1, \ldots, v_k\}$ of *macrovariables* ($k \in \mathbb{N}$). The v_i assume values in $\mathbb{V}_i$, and capital letters V_i denote the values of v_i. A particular set of values $\{V_1, \ldots, V_k\}$ defines a *macrostate* $M_{V_1,\ldots,V_k}$. We only write M rather than $M_{V_1,\ldots,V_k}$ if the specific V_i do not matter. A set of macrostates is complete if and only if it contains all states a system can be in.

A crucial posit of BSM is supervenience: a system's microstate uniquely determines its macrostate. Every macrostate M is associated with a *macroregion* Γ_M consisting of all $x \in \Gamma_E$ for which the system is in M. For a complete set of macrostates, the Γ_M form a partition of Γ_E (they do not overlap and jointly cover Γ_E).[2]

The *Boltzmann entropy* of a macrostate M is $S_B(M) := k_B \log[\mu(\Gamma_M)]$, where k_B is the Boltzmann constant. The Boltzmann entropy of a system at time t, $S_B(t)$, is the entropy of the system's macrostate at t: $S_B(t) := S_B(M_{x(t)})$, where $x(t)$ is the system's microstate at t and $M_{x(t)}$ is the macrostate supervening on $x(t)$.

We denote the equilibrium macrostate by M_{eq} and its macroregion by $\Gamma_{M_{\mathrm{eq}}}$. A crucial aspect of the standard presentation of BSM is that Γ_{eq} takes up most of Γ_{M_E}. To facilitate the discussion, we introduce the term β-dominance: $\Gamma_{M_{\mathrm{eq}}}$ is *β-dominant* if and only if $\mu(\Gamma_{M_{\mathrm{eq}}}) \geq \beta$ for $\beta \in (1/2, 1]$.

1. For details, see Frigg (2008, 103–21).

2. Because of a lack of space, our focus is on the most common case in which macrostates are defined relative to Γ_E. Our arguments generalize to cases in which the macrostates are defined relative to other subsets of Γ.

Often equilibrium is characterized as a state in which β is close to 1, but nothing in what follows depends on a particular choice of β.

The characterization of equilibrium as a β-dominant state goes back to Ehrenfest and Ehrenfest (1959, 30). While different versions of BSM explain the approach to equilibrium differently, β-dominance is a key factor in all of them. Those who favor an explanation based on ergodic theory have to assume that $\Gamma_{M_{eq}}$ takes up the majority of Γ_E because otherwise the system would not spend most of the time in $\Gamma_{M_{eq}}$ (e.g., Frigg and Werndl 2011, 2012). Those who see the approach to equilibrium as the result of some sort of probabilistic dynamics assume that $\Gamma_{M_{eq}}$ takes up most of Γ_E because they assign probabilities to macrostates that are proportional to $\mu(\Gamma_M)$ and that equilibrium comes out as the most likely state only if the equilibrium macroregion is β-dominant (e.g., Boltzmann 1877). Proponents of the typicality approach see dominance as the key ingredient in explaining the approach to equilibrium and sometimes even seem to argue that systems approach equilibrium because the equilibrium region takes up nearly all of phase space (e.g., Goldstein and Lebowitz 2004). We do not aim to adjudicate between these different approaches. Our question is a more basic one: Why is the equilibrium state β-dominant?

3. Justificatory Strategies. A look at the literature reveals three justificatory strategies. In practice these are often pursued side by side and seen as providing mutual support to each other. We assess each of them and argue that none of them is conclusive.

3.1. The Largest Number of Microstates and the Combinatorial Argument. The leading idea of the first justificatory strategy is that equilibrium is the macrostate that is compatible with the largest number of microstates. This strategy is exemplified by Boltzmann's (1877) combinatorial argument.[3] The state of one particle is determined by a point in its six-dimensional state space Γ_μ, and the state of a system of n identical particles is determined by n points in this space. Since the system is confined to a finite container and has constant energy E, only a finite part of Γ_μ is accessible. Boltzmann partitions the accessible part of Γ_μ into cells of equal size $\delta\omega$ whose dividing lines run parallel to the position and momentum axes. The result is a finite partition $\Omega := \{\omega_1, \ldots, \omega_m\}$, $m \in \mathbb{N}$. The cell in which a particle's state lies is its *coarse-grained microstate*. The coarse-grained microstate of the entire gas, called an *arrangement*, is given by a specification of the coarse-grained microstate of each of particle.

The system's macroproperties depend only on how many particles there are in each cell and not on which particles these are. A specification of the 'oc-

3. For details, see Uffink (2007) and Frigg (2008).

cupation number' of each cell is known as a *distribution* $D = (n_1, n_2, \ldots, n_m)$, where n_i is the number of particles whose state is in cell ω_i. Since m and n are finite, there are only finitely many distributions $D_1, \ldots, D_k$. Each distribution is compatible with several arrangements, and the number $G(D)$ of arrangements compatible with a given distribution $D = (n_1, n_2, \ldots, n_m)$ is

$$G(D) = \frac{n!}{n_1!n_2! \ldots n_m!}. \tag{1}$$

Every microstate x of Γ_E is associated with exactly one distribution $D(x)$. One then defines the set Γ_D of all x that are associated with a distribution D:

$$\Gamma_D = \{x \in \Gamma_E : D(x) = D\}. \tag{2}$$

Since macroproperties are fixed by the distribution, distributions are associated with macrostates. So we ask: Which of the distributions is the equilibrium distribution? Now Boltzmann's main idea enters the scene: equilibrium is the macrostate that is compatible with the largest number of microstates. To determine the equilibrium distribution, Boltzmann assumed that the energy e_i of particle i depends only on the cell in which it is located. Then the total energy is

$$\sum_{i=1}^{m} n_i e_i = E. \tag{3}$$

He furthermore assumed that the number of cells in Ω is small compared to the number of particles (allowing him to use Stirling's formula). With the further trivial assumption that $\sum_{i=1}^{m} n_i = n$, Boltzmann shows that $\mu_E(\Gamma_D)$ is maximal when

$$n_i = \gamma e^{\lambda e_i}, \tag{4}$$

where γ and λ are parameters that depend on n and E. This is the *discrete version of the Maxwell-Boltzmann distribution*. Thus, the equilibrium macrostate corresponds to the Maxwell-Boltzmann distribution.[4]

Its ingenuity notwithstanding, the combinatorial argument faces a number of important problems. The first is that it only applies to systems of noninteracting particles (Uffink 2007, 976–77). It provides a reasonable ap-

4. What (4) gives us is the distribution with the largest number of microstates (for the Lebesgue measure) on the 6N-dimensional shell-like domain Γ_{ES} specified by the condition (3). It does not give us the macroregion of maximal size (i.e., the distribution with the largest measure μ_E on the $6n - 1$ dimensional Γ_E). As Ehrenfest and Ehrenfest (1959, 30) stress, the (not further justified) assumption is that the possible distributions and the proportion of the different distributions would not change if macrostates were instead defined on Γ_E.

proximation for systems with negligible interparticle forces, but any other system is beyond its scope. Statistical mechanics (SM) ought to be a general theory of matter, and so this is a serious limitation.

The second problem is the absence of a conceptual connection between equilibrium in thermodynamics (TD) and the idea that the equilibrium macrostate is the one that is compatible with the largest number of microstates. In TD equilibrium is defined as the state to which isolated systems converge when left to themselves and which they never leave once they have reached it. This has very little, if anything, in common with the kind of considerations underlying the combinatorial argument. This is a problem for anyone who sees BSM as a reductionist enterprise. While the precise contours of the reduction of TD to SM remain controversial, we are not aware of any contributors who maintain radical antireductionism. Thus, the disconnect between the two notions of equilibrium is a serious problem.

Two replies come to mind. The first points out that since $\Gamma_{M_{eq}}$ is the largest subset of Γ_E, systems approach equilibrium and spend most of their time in $\Gamma_{M_{eq}}$. This shows that the BSM definition of equilibrium is a good approximation to the TD definition. This is not true in general. Whether a system spends most of its time in the β-dominant $\Gamma_{M_{eq}}$ depends on the dynamics. If, for instance, the dynamics is the identity function, it is not true that a system out of equilibrium approaches equilibrium and spends most of its time there. The second reply points out that we know for independent reasons that the equilibrium distribution is the Maxwell-Boltzmann distribution. This argument will be discussed in the next subsection, and our conclusion will be guarded.

Finally, the combinatorial argument (even if successful) shows that the equilibrium macrostate is larger than any other macrostate. However, as Lavis (2005) points out, this need not imply that the equilibrium is β-dominant. There may be a large number of smaller macrostates that jointly take up a large part of Γ_E. So the combinatorial argument does in fact not show that equilibrium is β-dominant.

3.2. The Maxwell-Boltzmann Distribution. According to the next justificatory strategy, a system is in equilibrium when its particles approximately satisfy the Maxwell-Boltzmann distribution (eq. [4]; e.g., Penrose 1989). This approach is misguided because the Maxwell-Boltzmann distribution is in fact the equilibrium distribution for a limited class of systems only, namely, for systems consisting of particles with negligible interparticle forces. For particles with nonnegligible interactions, different distributions correspond to equilibrium (Gupta 2003). Furthermore, for many simple models such as the Ising model (Baxter 1982) or the Kac ring (Lavis 2008) the equilibrium macrostate also does not correspond to the Maxwell-Boltzmann distribution.

This is no surprise given that the two common derivations of the distribution in effect assume that particles are noninteracting. Boltzmann's (1877) derivation is based on equation (3), the assumption that the total energy is the sum of the energy of the individual particles. This is true only if the particles are noninteracting (i.e., for ideal gases). While many expect that the argument also goes through for dilute gases (where this assumption holds approximately), the argument fails for nonnegligible interactions. In Maxwell's 1860 derivation (see Uffink 2007) the noninteraction assumption enters via the postulate that the probability distributions in different spatial directions can be factorized, which is true only if there is no interaction between particles. For these reasons, the Maxwell-Boltzmann distribution is the equilibrium distribution for a limited class of systems only and cannot be taken as a general definition of equilibrium.

3.3. Maximum Entropy. A third strategy justifies dominance by maximum entropy considerations along the following lines:[5] we know from TD that, if left to itself, a system approaches equilibrium, and equilibrium is a maximum entropy state. Hence, the Boltzmann entropy of a macrostate S_B is maximal in equilibrium. Since S_B is a monotonic function, the macrostate with the largest Boltzmann entropy is also the largest macrostate, which is the desired conclusion.

There are serious problems with the understanding of TD in this argument as well as with its implicit reductive claims. First, that a system, when left to itself, reaches equilibrium where entropy is maximal is often taken to be a consequence of the second law of TD, but it is not. As Brown and Uffink (2001) pointed out, that systems tend to approach equilibrium has to be added as an independent postulate, which they call the Minus First Law. But even if TD is amended with the Minus First Law, the conclusion does not follow. TD does not attribute an entropy to systems out of equilibrium. Thus, characterizing the approach to equilibrium as a process of entropy increase is meaningless from a TD point of view.

Even if all these issues could be resolved, there still would be a question why the fact that the TD entropy reaches a maximum in equilibrium would imply that the same holds for the Boltzmann entropy. To justify this inference, one would have to assume that the TD entropy reduces to the Boltzmann entropy. But it is far from clear that this is so. A connection between the TD entropy and the Boltzmann entropy has been established for ideal gases only, where the Sackur-Tatrode formula can be derived from BSM,

5. This strategy has been mentioned to us in conversation but it is hard to track down in print, at least in pure form. Albert's (2000) considerations concerning entropy seem to gesture in the direction of the third strategy.

which shows that both entropies have the same functional dependence on TD state variables. No such results are known for systems with interactions. Furthermore, there are well-known differences between the TD entropy and the Boltzmann entropy. Most importantly, the TD entropy is extensive while the Boltzmann entropy is not (Ainsworth 2012). But an extensive concept cannot reduce to a nonextensive concept (at least not without further qualifications). For these reasons, we conclude that maximum entropy considerations cannot be used to argue for the β-dominance of the equilibrium state.

4. Rethinking Equilibrium. The failure of standard justificatory strategies prompts the search for an alternative answer. In this section, we propose an alternative definition of equilibrium and introduce a new mathematical theorem proving that the equilibrium state thus defined is β-dominant.

The above strategies run into difficulties because there is no clear connection between the TD definition of equilibrium and β-dominance. Our aim is to provide the missing connection by taking as a point of departure the standard TD definition of equilibrium and then exploiting supervenience to 'translate' this macrodefinition into microlanguage.

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). In more detail: equilibrium is the state to which an isolated system converges when left on its own and which it never leaves once it has been reached (Callender 2001; Uffink 2001). Equilibrium in TD is unique in the sense that the system always converges toward the same equilibrium state. This leads to the following definition (the qualification 'strict' will become clear later):

> *Definition 1: Strict BSM Equilibrium*. Consider an isolated system S whose macrostates are specified in terms of the macrovariables $\{v_1, \ldots, v_k\}$, described as the measure-preserving dynamical system $(\Gamma_E, \Sigma_E, \mu_E, \phi_t)$. Let $M(x)$ be the macrostate that supervenes on microstate $x \in \Gamma$. Let $\Gamma_{M_{V_1,\ldots,V_k}} := \{x \in \Gamma_E : M(x) = M_{V_1,\ldots,V_k}\}$ be the set of all microstates on which $M_{V_1,\ldots,V_k}$ supervenes. If there is a macrostate $M_{V_1^*,\ldots,V_k^*}$ satisfying the following condition, then it is the strict BSM equilibrium state of S: for all initial states $x \in \Gamma_E$ at t_0 there exists a time t^* such that $M_{V_1,\ldots,V_k}(\phi_t(x)) = M_{V_1^*,\ldots,V_k^*}$ for all $t \geq t^*$.[6] We then write $M_{\text{eq}} := M_{V_1^*,\ldots,V_k^*}$.

Note that this definition incorporates the Minus First Law of TD.

6. If one wants to avoid the t^*-dependence on the initial state, one can instead demand that there exists a time t^* such that $v_i(t) = V_i^*$ for all initial states $M_{V_1,\ldots,V_k}$ and all $t \geq t^*$, $i = 1, \ldots, k$.

Before reflecting on this definition, we want to add a brief comment about reductionism. Reductive eliminativists may feel that a definition of equilibrium in SM that is based on 'top down translation' of its namesake in TD undermines the prospect of a reduction of TD to SM. They would argue that equilibrium has to be defined in purely mechanical terms and must then be shown to line up with the TD definition of equilibrium.

This point of view is not the only game in town, and reduction can be had even if equilibrium is defined 'top down' (as in the above definition).[7] First, whether the above definition undercuts a reduction depends on one's concept of reduction. For someone with a broadly Nagelian perspective, there is no problem: the above definition provides a bridge law, which allows the derivation of the requisite macroregularities from the laws of the microtheory. And a similar argument can be made in the framework of New Wave Reductionism. Second, equilibrium is a macroconcept: when describing a system as being in equilibrium, we look at it in terms of macroproperties. From a micro point of view there are only molecules bouncing around. They always bounce—there is no such thing as a relaxation of particle motion to an immutable state. Hence, a definition of equilibrium in macroterms is no heresy.

Definition 1 is too rigid for two reasons. The first reason is Poincaré recurrence: as long as the 'M' in SM refers to a mechanical theory that conserves phase volume (and there is widespread consensus that this is the case),[8] any attempt to justify an approach to strict equilibrium in mechanical terms is doomed to failure. The system will at some point return arbitrarily close to its initial condition, violating strict equilibrium (Uffink 2007; Frigg 2008). The second reason is that such a justification is not only unattainable but also undesirable. Experimental results show that equilibrium is not the immutable state that classical TD presents us with because systems exhibit fluctuations away from equilibrium (Wang et al. 2002). Thus, strict equilibrium is actually unphysical.

Consequently, strict definitions of equilibrium are undesirable both for theoretical and experimental reasons. So let us relax the condition that a system has to remain in equilibrium for all $t \geq t^*$ by the weaker condition that it has to be in equilibrium most of the time:

Definition 2: BSM α-Equilibrium. Consider the same system as in definition 1. Let $f_{M,x}(t)$ be the fraction of time of the interval $[t_0, t_0 + t]$ in which the system's state is in M when starting in initial state x at t_0, and let α be a real number in $[0.5, 1]$. If there is a macrostate $M_{V_1^*, \ldots, V_k^*}$ satisfying the follow-

7. For a discussion, see Dizadji-Bahmani, Frigg, and Hartmann (2010).

8. Hamiltonian Mechanics falls within this class, but the class is much wider.

ing condition, then it is the α-equilibrium state of S: for all initial states $x \in \Gamma_E$, $f_{M_{V_1^*,\ldots,V_k^*},x}(t) \geq \alpha$ in the limit $t \to \infty$. We then write $M_{\alpha-\mathrm{eq}} := M_{V_1^*,\ldots,V_k^*}$.

An obvious question concerns the value of α. Often the assumption seems to be that α is close to 1. This is reasonable but not the only possible choice. For our purposes nothing hangs on a particular choice of α, and so we leave it open what the best choice would be.

One last step is needed to arrive at the definition of equilibrium suitable for BSM. It has been pointed out variously that in SM, unlike in TD, we should not expect every initial condition to approach equilibrium (see, e.g., Callender 2001). Indeed, it is reasonable to allow for a set of very small measure ε for which the system does not approach equilibrium:

Definition 3: BSM α-ε-Equilibrium. Let S and $f_{M,x}(t)$ be as above. Let α be a real number in $[0.5, 1]$, let $1 > \varepsilon \geq 0$ be a small real number, and let Y be a subset of Γ_E such that $\mu_E(Y) \geq 1 - \varepsilon$. If there is a macrostate $M_{V_1^*,\ldots,V_k^*}$ satisfying the following condition, then it is the α-ε-equilibrium state of S: for all initial states $x \in Y$, $f_{M_{V_1^*,\ldots,V_k^*},x}(t) \geq \alpha$ in the limit $t \to \infty$. We then write $M_{\alpha-\varepsilon-\mathrm{eq}} := M_{V_1^*,\ldots,V_k^*}$.

Let us introduce the characteristic function of Γ_M, $1_M(x)$: $1_M(x) = 1$ for $x \in \Gamma_M$ and 0 otherwise. Definition 3 implies that, for all $x \in Y$,[9]

$$\lim_{T\to\infty} \frac{1}{T} \int_0^T 1_{M_{\alpha-\varepsilon-\mathrm{eq}}}(\phi_t(x))\, dt \geq \alpha. \tag{5}$$

An important assumption in this characterization of equilibrium is that μ_E (and not some other measure) is the relevant measure. It is often argued that μ_E can be interpreted as a probability or typicality measure (Frigg and Hoefer 2010; Werndl 2013). The condition then says that the system's state spends more than a fraction α of its time in equilibrium with probability $1 - \varepsilon$, or that typical initial conditions lie on trajectories that spend more than α of their time in equilibrium.

We contend that the relevant notion of equilibrium in BSM is α-ε-equilibrium. The central question then becomes: Why is the α-ε-equilibrium state β-dominant? Definition 3 in no way prejudges this question: it says nothing about the size of $\Gamma_{M_{\alpha-\varepsilon-\mathrm{eq}}}$, nor does it in an obvious sense imply anything about it.

That $\Gamma_{M_{\alpha-\varepsilon-\mathrm{eq}}}$ is β-prevalent follows from the following theorem, which we prove in the appendix:

9. This shows that definition 4 is closely related to Lavis's (2005, 255) characterization of TD-likeness.

Equilibrium Theorem: If $\Gamma_{M_{\alpha-\varepsilon-\mathrm{eq}}}$ is an α-ε-equilibrium of system S, then $\mu(\Gamma_{M_{\alpha-\varepsilon-\mathrm{eq}}}) \geq \alpha(1-\varepsilon)$.

We emphasize that the theorem is completely general in that no dynamical assumption is made (in particular, it is not assumed that the system is ergodic). So the theorem also applies to strongly interacting systems such as solids and liquids.

The equilibrium theorem is the centerpiece of our account. It shows in full generality that if the system S has an α-ε-equilibrium, then the equilibrium state is β-dominant for $\beta \geq \alpha(1-\varepsilon)$.[10] This provides the sought-after justification of the β-dominance of the equilibrium state.

The equilibrium theorem makes the conditional claim that if there is an α-ε-equilibrium, then $\mu(\Gamma_{M_{\alpha-\varepsilon-\mathrm{eq}}}) \geq \alpha(1-\varepsilon)$. As with all conditionals, the crucial and often vexing question is whether, and under what conditions, the antecedent holds. Some systems do not have equilibria. For instance, if the dynamics is given by the identity function, then no approach to equilibrium takes place, and the antecedent of the conditional is wrong. By contrast, epsilon ergodicity allows for an equilibrium state to exist (Frigg and Werndl 2011). This raises the question under which circumstances the antecedent is true, which is an important question for future research.

5. Conclusion. BSM partitions the phase space of a system into cells of macroscopically indistinguishable microstates. These cells are associated with the system's macrostates, and the largest cell is identified with equilibrium. What justifies the association of equilibrium with the largest cell? We discussed three justificatory strategies that can be found in the literature: that equilibrium is the macrostate compatible with the largest number of microstates, that equilibrium corresponds to the Maxwell-Boltzmann distribution and that most states are characterized by that distribution, and that equilibrium is the maximum entropy state. We argued that none of them is successful. This prompted the search for an alternative answer. We characterized equilibrium as the state in which the system spends most of its time and presented a new mathematical theorem proving that such an equilibrium state indeed corresponds to the largest cell. This result is completely general in that it is not based on any assumptions about either the system's dynamics or the nature of interactions within the system. It therefore provides the first fully general justification of the claim that the equilibrium state takes up most of the accessible part of the system's phase space.

10. It is assumed that ε is small enough so that $\alpha(1-\varepsilon) \geq 0.5$.

Appendix

Proof of the Equilibrium Theorem

The proof appeals to the ergodic decomposition theorem (cf. Petersen 1983, 81), stating that for a dynamical system $(\Gamma_E, \Sigma_E, \mu_E, \phi_t)$ the set Γ_E is the disjoint union of sets X_ω, each equipped with a σ-algebra Σ_{X_ω} and a probability measure μ_ω, and ϕ_t acts ergodically on each $(X_\omega, \Sigma_{X_\omega}, \mu_\omega)$. The indexing set is also a probability space $(\Omega, \Sigma_\Omega, P)$, and for any square integrable function f it holds that

$$\int_{\Gamma_E} f\, d\mu_E = \int_\Omega \int_{X_\omega} f\, d\mu_\omega\, dP. \tag{A1}$$

Application of the ergodic decomposition theorem for $f = 1_{M_{\alpha-\varepsilon-\mathrm{eq}}}(x)$ yields

$$\mu_E(\Gamma_{M_{\alpha-\varepsilon-\mathrm{eq}}}) = \int_{\Gamma_E} 1_{M_{\alpha-\varepsilon-\mathrm{eq}}}(x)\, d\mu_E = \int_\Omega \int_{X_\omega} 1_{M_{\alpha-\varepsilon-\mathrm{eq}}}(x)\, d\mu_\omega\, dP. \tag{A2}$$

For an ergodic system $(X_\omega, \Sigma_{X_\omega}, \mu_\omega, \phi_t)$, the long-run time average equals the phase average. Hence, for almost all $x \in X_\omega$

$$\lim_{T\to\infty} \frac{1}{T} \int_0^T 1_{M_{\alpha-\varepsilon-\mathrm{eq}}}(\phi_t(x))\, dt = \int_{X_\omega} 1_{M_{\alpha-\varepsilon-\mathrm{eq}}}(x)\, d\mu_\omega = \mu_\omega(\Gamma_{M_{\alpha-\varepsilon-\mathrm{eq}}} \cap X_\omega). \tag{A3}$$

From requirement (5) and because ϕ_t acts ergodically on each $(X_\omega, \Sigma_{X_\omega}, \mu_\omega)$, for almost all $x \in X_\omega$, $X_\omega \subseteq Y$

$$\alpha \le \lim_{T\to\infty} \frac{1}{T} \int_0^T 1_{M_{\alpha-\varepsilon-\mathrm{eq}}}(\phi_t(x))\, dt = \int_{X_\omega} 1_{M_{\alpha-\varepsilon-\mathrm{eq}}}(x)\, d\mu_\omega. \tag{A4}$$

Let us first consider the case $\varepsilon = 0$; that is, $\mu_E(Y) = 1$. Here, from equation (A2)

$$\mu_E(\Gamma_{M_{\alpha-\varepsilon-\mathrm{eq}}}) \ge \int_\Omega \alpha\, dP = \alpha. \tag{A5}$$

Hence, if $\varepsilon \ge 0$, it follows from equation (A2) that

$$\mu_E(\Gamma_{M_{\alpha-\varepsilon-\mathrm{eq}}}) \ge \alpha(1 - \varepsilon). \tag{A6}$$

REFERENCES

Ainsworth, Peter Mark. 2012. “Entropy in Statistical Mechanics.” *Philosophy of Science* 79: 542–60.

Albert, David. 2000. *Time and Chance*. Cambridge, MA: Harvard University Press.

Baxter, Rodney. 1982. *Exactly Solved Models in Statistical Mechanics*. San Diego, CA: Academic Press.

Boltzmann, Ludwig. 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.

Brown, Harvey, and Jos Uffink. 2001. “The Origins of Time-Asymmetry in Thermodynamics: The Minus First Law.” *Studies in History and Philosophy of Modern Physics* 32:525–38.

Callender, Craig. 2001. “Taking Thermodynamics Too Seriously.” *Studies in History and Philosophy of Modern Physics* 32:539–53.

Dizadji-Bahmani, Foad, Roman Frigg, and Stephan Hartmann. 2010. “Who Is Afraid of Nagelian Reduction?” *Erkenntnis* 73:393–412.

Ehrenfest, Paul, and Tatiana Ehrenfest. 1959. *The Conceptual Foundations of the Statistical Approach in Mechanics*. Ithaca, NY: Cornell University Press.

Frigg, Roman. 2008. “A Field Guide to Recent Work on the Foundations of Statistical Mechanics.” In *The Ashgate Companion to Contemporary Philosophy of Physics*, ed. Dean Rickles, 99–196. London: Ashgate.

Frigg, Roman, and Carl Hoefer. 2010. “Determinism and Chance from a Humean Perspective.” In *The Present Situation in the Philosophy of Science*, ed. Dennis Dieks, Wencelao Gonzales, Stephan Hartmann, Marcel Weber, Friedrich Stadler, and Thomas Übel, 351–72. Berlin: Springer.

Frigg, Roman, and Charlotte Werndl. 2011. “Explaining Thermodynamic-Like Behavior in Terms of Epsilon-Ergodicity.” *Philosophy of Science* 78:628–52.

———. 2012. “Demystifying Typicality.” *Philosophy of Science* 79:917–29.

Goldstein, Sheldon, and Joel L. Lebowitz. 2004. “On the (Boltzmann) Entropy of Nonequilibrium Systems.” *Physica* D 193:53–66.

Gupta, Mool C. 2003. *Statistical Thermodynamics*. New Delhi: New Age.

Lavis, David. 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.

Penrose, Roger. 1989. *The Emperor’s New Mind*. Oxford: Oxford University Press.

Petersen, Karl. 1983. *Ergodic Theory*. Cambridge: Cambridge University Press.

Reiss, Howard. 1996. *Methods of Thermodynamics*. Mineola, NY: Dover.

Uffink, Jos. 2001. “Bluff Your Way in the Second Law of Thermodynamics.” *Studies in History and Philosophy of Modern Physics* 32:305–94.

———. 2007. “Compendium of the Foundations of Classical Statistical Physics.” In *Philosophy of Physics*, ed. Jeremy Buttereld and John Earman, 923–1047. Amsterdam: North-Holland.

Wang, Genmiao, Edith M. Sevinck, Emil Mittag, Debra J. Searles, and Denis J. Evans. 2002. “Experimental Demonstration of Violations of the Second Law of Thermodynamics for Small Systems and Short Time Scales.” *Physical Review Letters* 89:050601.

Werndl, Charlotte. 2013. “Justifying Typicality Measures in Boltzmannian Statistical Mechanics.” *Studies in History and Philosophy of Modern Physics* 44:470–79.