

Equilibrium in Gibbsian 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	Justifying Equilibrium Distributions	4
3	Why Does Gibbsian Ensemble Averaging Work?	7
4	Gibbsian Non-Equilibrium Theory	10
5	Conclusion	14

1 Introduction

Consider a macroscopic system like a gas or a solid, and assume that this system is described by macroscopic variables such as temperature, pressure, volume and magnetisation. Intuitively a system has reached equilibrium when all change has come to a halt and the values of the macroscopic variables remain constant over time. How can equilibrium be characterised exactly, and why and how do systems that are not initially in an equilibrium state approach equilibrium? It is the aim of statistical mechanics (SM) to answer these questions in terms of the mechanical properties of the micro constituents of these systems and the dynamical laws that govern their time evolution.

Gibbsian SM (GSM) offers answers to these questions by associating an *ensemble* with each physical system. Let X be the state space of a system of interest. In a mechanical n -particle system, for instance, the state space has $6n$ dimensions, three dimensions for the position of each particle and three dimensions for the corresponding momenta. An ensemble is specified by a probability density $\rho(x, t)$ over the state space of a system, where t is time (which can be either continuous or discrete) and $x \in X$. Physical observables are associated with real-valued functions $f : X \rightarrow \mathbb{R}$. These functions represent physical quantities like internal energy, magnetisation, and polarisation. The *ensemble average* (or sometimes ‘phase average’) of an observable is:

$$\langle f \rangle = \int_X f(x) \rho(x, t) dx \quad (1)$$

Gibbs refers to the fact that ‘the distribution [...] will remain unchanged’ as the condition of ‘statistical equilibrium’ (1981 [1902]: 6). In modern parlance this amounts to saying that for an ensemble in equilibrium the distribution has to be stationary, meaning that it does not depend on time: $\rho(x, t) = \rho(x)$. Hence, equilibrium in GSM is the property of an ensemble rather than of an individual system. However, experimental tests are carried out on an individual physical system and not on an ensemble, which raises the question of what one can infer from an ensemble about such a system. The Gibbsian answer is that ensemble averages correspond to empirical equilibrium values: when measuring a physical quantity represented by f in a system in equilibrium, then the model tells us that the observed value will be $\langle f \rangle$, which does not depend on time because ρ is stationary.¹ This method is also known as *Gibbsian Ensemble Averaging* (or ‘*Gibbsian Phase Averaging*’).

Of the infinitely many stationary distributions three enjoy a special status: the microcanonical, canonical and grandcanonical distributions. The so-called *microcanonical*

¹For a discussion of the Gibbsian formalism and its interpretations see Frigg and Werndl (MS).

distribution describes the equilibrium of an ensemble consisting of systems in which the number of particles and the energy are both constant, and entire probability mass is evenly distributed on the hyper-surface of constant energy in X . The so-called *canonical distribution* is the equilibrium distribution of an ensemble of systems with a constant number of particles but varying energy. If both particle number and energy can vary, the relevant distribution is the so-called *grand-canonical distribution*.²

GSM raises three questions about equilibrium. First, GSM introduces three special distributions and specifies which situations they describe. These distributions are immensely powerful and are used in many applications. But what makes them special? Why, say, choose the microcanonical distribution when describing an isolated system and not one of the infinitely many other stationary distributions over X ? We address this issue in Section 2. Second, why does Gibbsian ensemble averaging work? GSM posits that what we observe are phase averages. In practice this is a successful method, but why do ensemble averages coincide with the values found in measurements performed on an actual physical system in equilibrium? There is no obvious connection between the two and if Gibbsian phase averaging is to be more than a black-box technique, we require an explanation of the connection between empirical values and ensemble averages. We address this question in Section 3. Third, a typical example of a non-equilibrium process is the following: a gas that is confined to the left half of a container starts expanding 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. As formulated so far GSM is an equilibrium theory, which does not offer a description of the approach to equilibrium. As we shall see, there is in fact a serious obstacle to extending GSM to non-equilibrium processes within the theory itself. The question is how this obstacle can be overcome. We discuss various suggestions in Section 4.

2 Justifying Equilibrium Distributions

There is an arbitrary number of stationary distributions over X and a further criterion is needed to single out a particular distribution as the one that corresponds to the equilibrium of an ensemble in a certain situation. A first suggestion is to select equilibrium distributions on grounds of simplicity (Phillies 2000: 112). But this criterion does not get off the ground as long as there is no hard and fast definition of simplicity, and even if there was such a definition it would remain unclear why equilibrium should correspond to a simple distribution.

²For formal definitions of these distributions see, for instance, Lavis and Bell (1999).

The *principle of indifference* says that as long as one has no reason to prefer one outcome over the other, one should not arbitrarily favour one outcome and instead assign equal probabilities to both. Tolman (1979 [1938]: 59-70) put this idea to use in SM when he introduced the principle that one should assign equal probabilities for a system to be in different parts of the state space when these parts have equal size (with respect to a relevant background measure on X). The microcanonical ensemble can then be seen as straightforward result of the application of this principle. However, the principle of indifference suffers from a number of well-known problems (for a discussion see Salmon *et al.* 1992: 74-77), and these problems also beset its application in GSM.

An important way to improve on the principle of indifference is an appeal to maximum entropy considerations. The Gibbs entropy of a distribution is

$$S_G[\rho] = -k \int_X \rho(x, t) \log[\rho(x, t)] dx, \quad (2)$$

where k is the Boltzmann constant. The *maximum entropy principle* then says that the correct distribution for certain situation is the one for which S_G is maximal with respect to the constraints imposed on the system in that situation (Kardar 2007: 34). The last clause is essential because different constraints lead to different equilibrium distributions. If, for instance, one keeps both the energy and the particle number in the system fixed, one can show that S_G is maximal for the microcanonical distribution. The special status of certain distributions in GSM are then justified by pointing out that they are the ones that satisfy the maximum entropy principle.

This pushes the justificatory question one step back: why should we accept the maximum entropy principle? Jaynes (1983) argues that the probabilities in statistical mechanics are subjective probabilities that represents our state of knowledge about the system (rather than any objective feature of the system itself). This allows him to appeal to Shannon's (1949) information theory to interpret entropy as a measure of the lack of information. A distribution with maximum entropy is then justified because it is maximally non-committal with respect to the available information. However, this justification of the maximum entropy principle remains controversial, both in GSM and in statistics more generally. Uffink (1995, 1996a) points out that the claim a maximisation procedure returns a unique distribution rests on unreasonably strong assumptions, and Howson and Urbach (2006: 276-88) doubt that a maximum entropy distribution really is non-committal.

A different route is taken by Malament and Zabell (1980), who aim to justify the microcanonical distribution by appeal to the dynamics of the system. A system is

ergodic on X if and only if for all measurable functions on X and for almost all initial conditions $x \in X$ (that is, for all initial conditions except, perhaps, for ones that taken together for a set of measure zero) the infinite time average of the function is equal to the ensemble average.³ Intuitively this means that if a set $A \in X$ has a measure of, say, $1/3$, then almost all trajectories will spend $1/3$ of their time in A in the long run. The significance of ergodicity in the current context lies in the fact that there is a uniqueness theorem for measures in ergodic systems. A measure λ on X is *absolutely continuous* with respect to a measure μ on X iff whenever μ assigns measure zero to a set then λ also assigns measure zero to that set. The uniqueness theorem says that if a system is ergodic with respect to the measure μ and λ is a measure that is (i) invariant under the dynamics of the system and (ii) absolutely continuous with respect to μ , then $\mu = \lambda$. Intuitively the theorem says that the ergodic measure is the only invariant measure in the class of all measures that are absolutely continuous with the ergodic measure. Malament and Zabell postulate that any acceptable measure has to be absolutely continuous with the ergodic measure. This is because if a subset of X is just a small displacement of another subset of X , then probability of finding the micro-state of the system in one set should be close to that of finding it in the other, which, as they show, is equivalent to absolute continuity. This singles the microcanonical measure as the only correct measure.

This approach is premised on systems being ergodic, but many systems which are dealt with successfully by GSM fail to meet this condition. For instance, in a solid the molecules oscillate around fixed positions in a lattice and therefore can only access a small part of the state space (Uffink 2007: 1017). Bricmont (2001) points out that neither the Kac Ring (Kac 1959) nor a system of n uncoupled anharmonic oscillators of identical mass is ergodic and yet both exhibit equilibrium behaviour. And, most notably, a system of non-interacting point particles is known not to be ergodic and many applications of SM are based on such systems (Uffink 1996b: 381). This is a problem because, as Earman and Rédei (1996: 71) point out, even if a system is just a ‘little bit’ non-ergodic, then the uniqueness theorem fails completely, and with it Malament and Zabell’s approach.

To circumvent this difficulty, Vranas (1998) suggested replacing ergodicity with what he calls ε -ergodicity. Intuitively, a system is ε -ergodic when it is ergodic on the vast majority of X , namely on a subset Y of X so that $\mu(Y) \geq 1 - \varepsilon$, where ε is a very small real number or zero (and it is assumed that μ is normalised). Vranas points out that there is a middle ground between ergodicity holding and failing completely. He articulates a notion of two measures being ε -close and proves a ε version of the uniqueness theorem. This allows him to reformulate Malament and Zabell’s approach

³See Berkovitz *et al.* (2011) for a non-technical introduction to ergodic theory.

for ε -ergodic systems. This is progress because the class of systems that are ε -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 ε -ergodic and it remains unclear how the current approach can deal with them. And, most notably, even if all these difficulties could be resolved, the approach only concerns the microcanonical distribution and it remains unclear how the canonical and the grand-canonical distributions can be justified.

Finally, Wallace (2001) argues that there is no reason to think that the classical definitions and derivations of equilibrium apply to quantum mechanical systems. Because of this, he argues, these definitions and derivations cannot be valid even in the classical regime. If they can be made to work for classical systems, then this must be an artefact of classical mechanics that has nothing to do with the actual world. Therefore, he argues, a new notion of quantum mechanical equilibrium is needed. Wallace does not articulate such a new notion and it remains an open question how the new quantum mechanical equilibrium would look like. Finding a quantum mechanical justification of probabilities in statistical mechanics is worthwhile endeavour. Yet it remains unclear whether far-reaching conclusions about classical justification can be drawn. In particular, it is not clear why the fact that classical derivations of the equilibrium distributions do not work in quantum mechanics automatically implies that the justifications for the probability distributions in the classical regime also fail. Even if at the fundamental level the world is quantum, there can be different justifications for different levels of reality.

3 Why Does Gibbsian Ensemble Averaging Work?

Why does Gibbsian ensemble averaging work? That is, why do phase averages coincide with values measured in actual physical systems? Common wisdom justifies the use of phase averages by appeal to time averages.⁴ GSM associates a physical quantity with a real-valued function f on X . Performing a measurement takes time, and so a measurement device eventually registers a time average of f . Although making a measurement only takes a short time by human standards, the time is long compared to the time that a typical molecular process takes. For this reason, the measured value is approximately equal to the *infinite* time average of f . If one now assumes that the system is ergodic, time and ensemble averages are equal. This, so the argument concludes, shows that the ensemble averages one can calculate with

⁴This view is discussed but not endorsed, for instance, in Bricmont (1996: 145-46), Earman and Rédei (1996: 67-69), Malament and Zabell (1980: 342), and van Lith (2001a: 581-83).

GSM are the values obtained from measurements.

This argument faces a number of difficulties. First, from the fact that measurements take time it does not follow that the measurement outcome is a time average. Measurement devices could just as well record an instantaneous value at some point during the measurement. Second, even if it was the case that measurement outcomes reflected time averages, equating a *finite* time average with an *infinite* time average is problematic. Even if a measurement takes a long time compared to molecular processes, finite and infinite averages could be very different (Malament and Zabell 1980: 342-43; Sklar 1973: 211). This gives cause for concern because ergodicity *only* holds for infinite time averages. So if infinite time averages are replaced by finite time averages one can no longer appeal to ergodicity to equate time and ensemble averages, and this pulls the rug from underneath the above argument. Third, as we have seen above, many systems which are dealt with successfully by GSM are not ergodic on X .

Khinchin (1960 [1949]) addresses the third problem by pointing out that full-fledged ergodicity is unnecessary. He sees the key to an understanding of GSM in a careful study of the the class of functions that are relevant for SM. He submits that there are two restrictions. First, only systems with a large number of degrees of freedom are considered because SM studies systems that are made up of a large number of micro-constituents. Second, attention should be restricted to so-called *sum functions*, i.e. functions that are a sum over one-particle functions. Restricting attention to Hamiltonian systems and assuming that the Hamiltonian is itself a sum function, Khinchin could prove that for every sum function f there are constants k_1 and k_2 such that the measure of all states in X for which $|(f^*(x) - \langle f \rangle) / \langle f \rangle|$ is greater than $k_1 n^{-1/4}$ is smaller than $k_2 n^{-1/4}$, where $f^*(x)$ is the time average of f on a trajectory starting in x . This theorem is sometimes also referred to as ‘Khinchin’s ergodic theorem’.⁵ Intuitively, it says that as n becomes larger, the measure of those regions of X where the time and the space average differ by more than a small amount tends towards zero.

This approach faces a number of challenges. First, like the original ergodic approach it associates measurement outcomes with infinite time averages and is therefore vulnerable to the same objections. Second, the motivation for focussing attention on sum-functions is that this is the relevant class of functions for SM. Batterman (1998: 191) points out that this is too narrow as there are functions of interest that do not have this form, in particular the Hamiltonian of systems in which particles interact with each other. Third, and most importantly, there is what Khinchin himself called the ‘methodological paradox’ (Khinchin 1960 [1949]: 41-43). The proof of Khinchin’s

⁵For a discussion of the proof see Badino (2006), Batterman (1998: 190-98) and van Lith (2001b: 83-90).

theorem assumes that the Hamiltonian is a sum function, which is tantamount to assuming that the components of the system do not interact. However, non-interacting systems often do not reach equilibrium. If, for instance, the particles of a gas do not collide, the gas does not converge towards an equilibrium state. So Khinchin's theorem seems to describe the equilibrium of a class of systems that don't reach equilibrium! Khinchin suggested avoiding the problem by assuming that there are only short range interactions between the molecules. These are sufficient to bring the system into equilibrium and at the same time have no significant effect on averages.

This response seems *ad hoc* and a better solution is needed. This was the starting point for a research programme now known as the *thermodynamic limit* which aims to establish 'Khinchin-like' results for Hamiltonians *with* interaction terms. Results of this kind can be proven in the limit for $n \rightarrow \infty$, if also the volume V of the system tends towards infinity in such a way that the number-density n/V remains constant. Classic statements are Ruelle (1969, 2004); surveys and further references can be found in Compagner (1989), Luczak (2016) and Uffink (2007: 1020-28).

Malament and Zabell's approach (which we have introduced in Section 2) rejects the idea that we observe time averages and posits that if a system is in state $x \in X$ at time t and a measurement is performed at t , then the observed value is $f(x)$. We nevertheless observe phase averages if the function f exhibits small dispersion with respect to the microcanonical measure, meaning that the set of points in X where the value of f significantly differs from $\langle f \rangle$ is vanishingly small. If this is the case, then, at any given time, it is overwhelmingly likely that the micro-state of the system is one for which the value of f coincides with $\langle f \rangle$, which explains why Gibbsian ensemble averaging works. To justify that functions meet this criterion, Malament and Zabell refer to Khinchin's research programme. However, as we have seen, there are questions about Khinchin's assumptions. Furthermore, it is not the case that for all observables f relevant in practice, f exhibits small dispersion with respect to the microcanonical measure (cf. Werndl and Frigg forthcoming, MS). Finally, as noted above, Malament and Zabell's approach addresses only the microcanonical ensemble and even if all internal issues could be resolved, there would still be a question about the justification of the canonical and the grandcanonical distributions.

Werndl and Frigg (forthcoming, MS) suggest that the solution to the problem lies in a reassessment of the status of GSM. The tacit assumption shared by all accounts discussed so far is that GSM has to apply to all situations and give correct results in all cases. Werndl and Frigg argue that this not so. The true fundamental theory is Boltzmannian statistical mechanics (see chapter 7.1, on Equilibrium in Boltzmannian Statistical Mechanics), and GSM is an effective theory in that it offers a set of prin-

principles that make the calculation of outcomes possible in cases where the fundamental theory is too difficult to handle. Like all effective theories, GSM buys computational tractability at the cost of incompleteness, and for this reason it has a limited range of applicability. The crucial question therefore is under what circumstances Gibbsian SM yields trustworthy results and when its procedures fail to deliver. Werndl and Frigg state two sufficient (but not necessary) conditions for GSM to provide correct results. The first is that the relevant phase function has a small dispersion, i.e. assumes (almost) the same value almost everywhere; the second is that the relevant observable is the sum of one component observables defined on partition of the one particle space where all values of the observable have the same probability. This at once provides a justification of GSM where it is successful, and an explanation of why it fails when it does.

4 Gibbsian Non-Equilibrium Theory

A gas that is confined to the left half of a container starts expanding when the confining wall is removed and eventually fills the available space evenly. This is a standard example of a system approaching equilibrium when initially prepared in a non-equilibrium state. In GSM the distribution $\rho(x, t)$ specifies the state of the Gibbsian ensemble and so one would expect that the approach to equilibrium is described in terms of changes in $\rho(x, t)$. The idea would be that there would be a distribution associated with the gas confined to the left half (presumably peaked over the parts of X that correspond to the left half of the container and zero on the other parts of X), which then evolves toward the stationary equilibrium distribution of maximum entropy.

Large parts of GSM are carried out against the background of Hamiltonian Mechanics, meaning that the time evolution of the system is assumed to be governed by Hamilton's equation of motion. Unfortunately, Hamiltonian mechanics does not sit well with the expectations described in the last paragraph. First, the Gibbs entropy turns out to be a constant of motion (Zeh 2001: 48-49), meaning that it does not change over time. This precludes a characterisation of the approach to equilibrium in terms of increasing entropy. Second, the Hamiltonian equations of motion preclude the evolution from a non-stationary to a stationary distribution. If, at some point in time, the distribution is non-stationary, then it will remain non-stationary for all times and, conversely, if it is stationary at some time, then it must have been stationary all along (van Lith 2001a: 591-92). Hence either the system has either always been in equilibrium, or else it will never be. This is serious problem for characterisation of equilibrium in terms of stationary distributions because it implies that

systems never approach equilibrium.

This has sparked a number of attempts to remedy the situation. A first suggestion is based on the method of *coarse-graining*. The method was introduced by Gibbs (1981 [1902]: Ch. 12) and has since gained prominence in the literature on GSM (Penrose 1970). The idea is to put a grid on the state space X consisting of cells of finite size (hence the expression ‘coarse-graining’), and then define the coarse-grained density $\bar{\rho}$ as the density that is uniform within each cell and takes the average of ρ in the cell as its value. One can then plug $\bar{\rho}$ into Equation 2, which yields the coarse-grained entropy \bar{S}_G . The coarse-grained density is not subject to the same restrictions as the original density and so it is, at least in principle, possible for $\bar{\rho}$ to evolve in such a way that it will converge to maximum entropy distribution that does not change any more once it has been reached. This state is the *coarse-grained equilibrium* (Ridderbos 2002: 69). The approach to the coarse-grained equilibrium is accompanied by an increase in the coarse grained entropy, which, unlike the original Gibbs entropy, is not a constant of motion. The question then is under what circumstances $\bar{\rho}$ tends toward a coarse-grained equilibrium. The standard answer is that the dynamics of the system has to be *mixing*. Intuitively, a system is mixing if, in the long run, every subset A of X gets evenly scattered over the entire X .⁶ In sum, if one replaces ρ and S_G by $\bar{\rho}$ and \bar{S}_G , and assumes the the dynamics is mixing, then the system reaches a coarse-grained equilibrium.

This programme faces a number of challenges. First, mixing is a stringent requirement, which many systems fail to meet. In fact, mixing implies ergodicity and so if a system is not ergodic, it is not mixing either. We have seen above that ergodicity is not easy to come by. Second, the thrust of the argument in favour of coarse-graining is that one cannot empirically distinguish between ρ and $\bar{\rho}$ and so one is at liberty to use $\bar{\rho}$ rather than ρ if this resolves theoretical issues in the theory. Ridderbos and Redhead (1998) argue that this was too hasty a conclusion because there are in fact experimental contexts in which it is possible to empirically distinguish between ρ or $\bar{\rho}$. One such context is the so-called spin-echo experiment (Hahn 1950), which studies how spins process in a magnetic field.⁷ In the experiment a spin system is prepared in a non-equilibrium state, and this state turns out to be such that the system evolves into a coarse-grained equilibrium. If the system is then exposed to a radio pulse, the spins end up returning to their initial state, thereby moving out of the coarse-grained equilibrium. This, so the argument goes, shows that one can empirically distinguish between a real equilibrium and coarse-grained equilibrium,

⁶For a non-technical introduction to mixing see Berkovitz *et al.* (2011).

⁷For detailed discussions of spin-echo experiment see Frigg (2008: 159-64) and Sklar (1993: 219-22).

which invalidates the rationale for replacing ρ by $\bar{\rho}$. This criticism has not gone unchallenged. Ainsworth (2005) defends coarse graining on the grounds that empirical equivalence is not necessary, and Lavis (2004) argues that the behaviour of $\bar{\rho}$ and the coarse-grained entropy predicted by Ridderbos and Redhead is an artefact of the way in which they calculated these quantities.

Van Lith also proposes to change the definition of equilibrium, but in a different way than the coarse-graining approach. She points out that the stationarity of ρ is actually too strong a requirement because all that is needed from thermodynamic point of view is that the distribution be such that the phase average of every function in a physically relevant set of functions only fluctuates mildly around its average value (van Lith 1999: 114). A distribution can meet this requirement even if it is not stationary. As in the coarse-graining approach the question arises under what circumstances averages behave in the required way, and the answer is also that mixing is a sufficient condition. But as already noted, mixing is a stringent condition. While mixing is not needed for van Lith's approach to equilibrium to take place (because her equilibrium condition is weaker), she does not propose an alternative condition that is sufficient for the approach to equilibrium and show that the relevant systems in practice fulfill this condition. Furthermore, as van Lith herself points out (van Lith 1999: 115), the proposal does not contain a recipe to get the Gibbs entropy moving and hence an approach to equilibrium is not accompanied by a corresponding entropy increase.

Interventionism takes the opposite route: it leaves the definitions of equilibrium and entropy intact and changes the dynamics of the system. So far we assumed that the systems under consideration are isolated. Interventionism renounces this idea and submits that taking outside influences on the system into account offers a solution to the problem. Blatt (1959) suggests that this is the source of the problem. The perturbations from outside the system in fact serve as a kind of 'stirring mechanism' that drives the system into equilibrium. These outside influences are describable only in statistical term and hence the system is no longer governed by the deterministic Hamiltonian equations. This removes the obstacle to the distribution converging to an unchanging maximum entropy distribution, and the entropy can increase. Blatt (1959) and Ridderbos and Redhead (1998) look at models of outside perturbations and report that under realistic assumptions systems indeed do approach equilibrium.

An obvious objection is that we are always free to consider a larger system consisting of our 'original' system *and* its environment (for instance, we can consider a 'gas plus container' system). Because Hamilton's equation are a universal theory, they will govern that system, and so we are where we started. The interventionist can

then reply that we can consider the environment of the environment (for instance, the laboratory in which the container is located) and consider outside influences, and the critic can make the same move as before. This leads to a regress that only ends when we consider the entire universe as one big system. Whether the interventionist really has to throw in the towel at this point depends on the status one assigns to the mechanical background theory, here Hamiltonian mechanics. For those who see this Hamiltonian mechanics as a truly universal theory this is the end of interventionism. But this understanding of mechanics is not uncontroversial. If, like Cartwright (1999), one thinks that we cannot legitimately claim that laws apply universally (and extends this sceptical argument to mechanics), then the argument against interventionism does not get off the ground.

Another approach that changes the dynamics of the system is the *stochastic dynamics* programme, which replaces the (deterministic) Hamiltonian equations of motion by stochastic equations. Typically this is done by coarse-graining the phase space and then postulating a probabilistic law describing the transition from one cell of the partition to another one. This removes the restrictions of Hamiltonian mechanics and makes room for the distribution to spread and the entropy to increase. Classical statements of this approach are Mackey (1992), Penrose (1970), and Streater (1995). The main question facing this account is how the stochastic equations can be justified. Its probabilistic equations are often introduced as independent postulates on a case-by-case basis, and it is unclear whether they follow from a underlying fundamental law of motion.⁸

Brussels School (sometimes also ‘Brussels-Austin School’) bases its analysis on the view that the notion of a precise state of a system (represented by a point x in X) is meaningless in a system with sensitive dependence on initial conditions. For this reason deterministic laws specifying the time evolution of points should be replaced by an explicitly probabilistic description of the system in terms of open regions of the phase space. This programme, if successful, can be seen as providing the sought after justification for the stochastic dynamics programme mentioned in the last paragraph.⁹

Jaynes’ approach, which we encountered in Section 2, also offers a response to the issues with non-equilibrium processes. Since ρ is a codification of a scientist’s knowledge rather than a feature of the system itself, it is not bound by mechanical equations of motion. This can be exploited to change the distribution so that the entropy

⁸For discussions of the stochastic dynamics programme see Callender (1999: 358-64), Sklar (1993: Chs. 6 and 7), and Uffink (2007: 1038-63).

⁹For presentations and critical discussions of the ideas of the Brussels School see Batterman (1991), Bishop (2004), and Karakostas (1996).

increases and the distribution spreads. This is done by iteratively updating the distribution every time an observation is made, ensuring that the distribution satisfies the maximum entropy principle after every observation (see Sklar (1993: 255-57) for a description of the details of this process). However, Lavis and Milligan (1985) point out that this procedure does not necessarily lead to a *monotonic* entropy increase, as conformity with thermodynamics would require. Furthermore, if ρ is a reflection of our knowledge, then changes in ρ are reflections on changes in our knowledge. This, so a number of critics have argued, has the consequence that the approach to equilibrium depends on what we know. But, so the argument goes, this is absurd because the boiling of kettles or the spreading of gases is a consequence of how molecules behave and not of what we know about them (Goldstein 2001: 48; Redhead 1995: 27-28).

5 Conclusion

We have introduced the Gibbsian notion of equilibrium and reviewed a number of issues that arise in its justification and use. To conclude we would like to draw attention to two further issues. The first issue is the interpretation of the probabilities in GSM. Gibbs himself seems to have thought about probabilities in a frequentist way, which is problematic for a number of reasons (Frigg 2008: 153-54). An alternative interpretation are time averages, but these do not fare better (Dougherty 1993; von Plato 1981). Bayesian probabilities face the objections mentioned in Section 2. So the interpretation of GSM probabilities remains a challenge. Second, the relation between GSM and Boltzmannian statistical mechanics is poorly understood. Even though both frameworks are invoked in practical applications as well as in foundational discussions, it remains unclear how they, and in particular their notions of equilibrium, relate to each other. But having two incompatible notions of equilibrium at work in SM is unsatisfactory and a reflection on how the two approaches relate is necessary. Steps towards a better understanding of the relationship between the two frameworks are made in Lavis (2005) and Werndl and Frigg (forthcoming), but the problem deserves further attention.

References

- Ainsworth, P. (2005) "The spin-echo experiment and statistical mechanics," *Foundations of Physics Letters* 18: 621-35.
- Badion, M. (2006) "The foundational role of ergodic theory," *Foundations of Science* 11: 323-47.

- Batterman, R. W. (1991) "Randomness and probability in dynamical theories: On the proposals of the Prigogine School," *Philosophy of Science* 58: 241-63.
- (1998) "Why equilibrium statistical mechanics works: Universality and the renormalization group," *Philosophy of Science* 65: 183-208.
- 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/>.
- Bishop, R. (2004) "Nonequilibrium statistical mechanics Brussels-Austin Style," *Studies in History and Philosophy of Modern Physics* 35: 1-30.
- Blatt, J. (1959) "An alternative approach to the ergodic problem," *Progress in Theoretical Physics* 22: 745-55.
- 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.
- (2001) "Bayes, Boltzmann, and Bohm: Probabilities in physics," in J. Bricmont, G. Ghirardi, D. Dürr, F. Petruccione, M. C. Galavotti and N. Zanghì (eds), *Foundations and Perspectives* (Lecture Notes in Physics, Vol. 574), Berlin and Heidelberg: Springer, pp. 4-21.
- Callender, C. (1999) "Reducing thermodynamics to statistical mechanics: The case of entropy," *Journal of Philosophy* 96: 348-73.
- Cartwright, N. (1999) *The Dappled World. A Study of the Boundaries of Science*, Cambridge: Cambridge University Press.
- Compagner, A. (1989) "Thermodynamics as the continuum limit of statistical mechanics," *American Journal of Physics* 57: 106-17.
- Dougherty, J. P. (1993) "Explaining statistical mechanics," *Studies in History and Philosophy of Science* 24: 843-66.
- 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.
- 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. and Werndl, C. (2011a) “Explaining Thermodynamic-Like Behaviour in Terms of Epsilon-Ergodicity,” *Philosophy of Science* 78: 628-52.
- Frigg R. and Werndl, C. (MS): *Interpreting Gibbs*, Manuscript.
- Gibbs, J. W. (1981 [1902]) *Elementary Principles in Statistical Mechanics*, Woodbridge: Ox Bow Press.
- 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.
- Hahn, E. L. (1950) “Spin echoes,” *Physics Review* 80: 580-94.
- Howson, C. and Urbach, P. (2006) *Scientific Reasoning. The Bayesian Approach*, Chicago and La Salle/IL: Open Court.
- Jaynes, E. T. (1983) *Papers on Probability, Statistics, and Statistical Physics*, ed by R. D. Rosenkrantz, Dordrecht: Reidel.
- Kac, M. (1959) *Probability and Related Topics in Physical Science*, New York: Interscience Publishing.
- Karakostas, V. (1996) “On the Brussels Schools arrow of time in quantum theory,” *Philosophy of Science* 63: 374-400.
- Kardar, M. (2007) *Statistical Physics of Particles*, Cambridge: Cambridge University Press.
- Khinchin, A. I. (1960 [1949]) *Mathematical Foundations of Statistical Mechanics*, Mineola/NY: Dover Publications.
- Lavis, D. (2004) “The spin-echo system reconsidered,” *Foundations of Physics* 34: 669-88.
- (2005) “Boltzmann and Gibbs: An attempted reconciliation,” *Studies in History and Philosophy of Modern Physics* 36: 245-73.
- Lavis, D. and Bell, G. (1999) *Statistical Mechanics of Lattice Systems: Volume 1: Closed-Form and Exact Solutions*, Berlin: Springer.
- Lavis, D. and Milligan, P. (1985) “Essay Review of Jaynes Collected Papers,” *British Journal for the Philosophy of Science* 36: 193-210.

- Luczak J. (2016) “On How to Approach the Approach to Equilibrium,” *Philosophy of Science* 83: 393-411.
- Mackey, M. C. (1992) *Times Arrow: The Origins of Thermodynamic Behaviour*, Berlin: Springer.
- Malament, D. and Zabell, S. L. (1980) “Why Gibbs phase averages work,” *Philosophy of Science* 47: 339-49.
- Penrose, O. (1970) *Foundations of Statistical Mechanics*, Oxford: Oxford University Press.
- Phillies, G.D.J. (2000) *Elementary Lectures in Statistical Mechanics*, Berlin: Springer.
- Redhead, M. L. G. (1995) *From Physics to Metaphysics*, Cambridge: Cambridge University Press.
- Ridderbos, K. (2002) “The coarse-graining approach to statistical mechanics: How blissful is our ignorance?,” *Studies in History and Philosophy of Modern Physics* 33: 6-77.
- Ridderbos, T. M. and Redhead, M. L. G. (1998) “The Spin-echo experiments and the second law of thermodynamics,” *Foundations of Physics* 28: 1237-70.
- Ruelle, D. (1969) *Statistical Mechanics: Rigorous Results*, New York: Benjamin.
- (2004) *Thermodynamic Formalism: The Mathematical Structure of Equilibrium Statistical Mechanics*, Cambridge: Cambridge University Press.
- Salmon, M. *et al.* (1992) *Introduction to the philosophy of science*, Indianapolis and Cambridge: Hackett.
- Shannon, C. E. (1949) *The mathematical theory of communication*, in C. E. Shannon and W. Warren, *The Mathematical Theory of Communication*, Urbana, Chicago and London: University of Illinois Press.
- Sklar, L. (1973) “Statistical explanation and ergodic theory,” *Philosophy of Science* 40: 194-212.
- (1993) *Physics and Chance. Philosophical Issues in the Foundations of Statistical Mechanics*, Cambridge: Cambridge University Press.
- Streater, R. F. (1995) *Stochastic Dynamics: A Stochastic Approach to Nonequilibrium Thermodynamics*, London: Imperial College Press.

- Tolman, R. C. (1979 [1938]) *The Principles of Statistical Mechanics*, Mineola/New York: Dover.
- Uffink J. (1995) “Can the maximum entropy principle be explained as a consistency requirement?,” *Studies in History and Philosophy of Modern Physics* 26: 223-61.
- (1996a) “The constraint rule of the maximum entropy principle,” *Studies in History and Philosophy of Modern Physics* 27: 47-79.
- (1996b) “Nought but molecules in motion” (Review essay of Lawrence Sklar: *Physics and Chance*), *Studies in History and Philosophy of Modern Physics* 27: 373-87.
- (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.
- van Lith, J. (1999) “Reconsidering the concept of equilibrium in classical statistical mechanics,” *Philosophy of Science* 66 (Supplement): 107-18.
- van Lith, J. (2001a) “Ergodic theory, interpretations of probability and the foundations of statistical mechanics,” *Studies in History and Philosophy of Modern Physics* 37: 581-94.
- (2001b) *Stir in Sillness: A Study in the Foundations of Equilibrium Statistical Mechanics*, PhD Thesis, University of Utrecht, (available at <http://www.library.uu.nl/digiarchief/dip/diss/1957294/inhoud.htm>).
- von Plato, J. (1981) “Reductive relations in interpretations of probability,” *Synthese* 48: 61-75.
- Vranas, P. B. M. (1998) “Epsilon-ergodicity and the success of equilibrium statistical mechanics,” *Philosophy of Science* 65: 688-708.
- Wallace (2001): “Implications of quantum theory in the foundations of statistical mechanics”, Manuscript, available in the PhilSci Archive at <http://philsci-archive.pitt.edu/410/1/wallace.pdf>
- Werndl C. and Frigg, R. (MS) *Gibbs vs Boltzmann: Equilibrium*, Manuscript.
- Werndl C. and Frigg, R. (forthcoming) “Mind the Gap: Boltzmannian vs Gibbsian Equilibrium,” forthcoming in *Philosophy of Science*.
- Zeh, H. D. (2001) *The Physical Basis of the Direction of Time*, Berlin and New York: Springer