Typicality and the Approach to Equilibrium in Boltzmannian Statistical Mechanics

Roman Frigg†‡

An important contemporary version of Boltzmannian statistical mechanics explains the approach to equilibrium in terms of typicality. The problem with this approach is that it comes in different versions, which are, however, not recognized as such and not clearly distinguished. This article identifies three different versions of typicality-based explanations of thermodynamic-like behavior and evaluates their respective successes. The conclusion is that the first two are unsuccessful because they fail to take the system's dynamics into account. The third, however, is promising. I give a precise formulation of the proposal and present an argument in support of its central contention.

1. Introduction. Consider a gas confined to the left half of a container. Removing the dividing wall results in the gas spreading uniformly over the entire available space. It has approached equilibrium. Statistical mechanics (SM) aims to explain the approach to equilibrium in terms of the dynamical laws governing 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? And why does this happen invariably? That is, why do we never observe gases staying in the left half, even after the shutter has been removed?

An important contemporary version of the Boltzmannian approach to SM, originating in the work of Joel Lebowitz (1993a, 1993b), answers these questions in terms of the notion of *typicality*. Intuitively, something

†To contact the author, please write to: Department of Philosophy, Logic, and Scientific Method, London School of Economics, Houghton Street, London WC2A 2AE, England; e-mail: r.p.frigg@lse.ac.uk.

‡Special thanks goes to Charlotte Werndl for invaluable discussions about dynamical systems. I also would like to thank Craig Callender, Hartry Field, Shelly Goldstein, Stephan Hartmann, Carl Hoefer, Wolfgang Pietsch, and Nino Zanghì for valuable comments on earlier drafts and helpful discussions. Thanks to the audiences in Madrid, Oxford, Pittsburgh, and Jerusalem for stimulating discussions.

Philosophy of Science, 76 (December 2009) pp. 997–1008. 0031-8248/2009/7605-0013$10.00

is typical if it happens in the 'vast majority' of cases: typical lottery tickets are blanks, typical Olympic athletes are well trained, and in a typical series of 1,000 coin tosses the ratio of the number of heads and the number of tails is approximately one. The aim of a typicality-based approach to SM is to show that approaching equilibrium is the typical behavior of systems like gases.

This approach has grown increasingly popular in recent years (references are given below). The problem with it is that it comes in different versions, which are, however, not recognized as such, much less clearly distinguished. The aim of this article is to distinguish three different kinds of typicality-based explanations of the approach to equilibrium and evaluate their respective successes. My conclusion is that the first two are unsuccessful because they fail to take the system's dynamics into account. The third, however, is promising. I give a precise formulation of the proposal and present the outline of a proof of its central contention.

2. Classical Boltzmannian SM. Consider a system consisting of n classical particles with 3 degrees of freedom each. The state of this system is specified by a point x, also referred to as the system's microstate, in its $6n$-dimensional phase space Γ, which is endowed with the Lebesgue measure μ_L. The dynamics of the system is governed by Hamilton's equations of motion, which define a measure preserving flow ϕ_t on Γ, meaning that for all times t $\phi_t : \Gamma \rightarrow \Gamma$ is a one-to-one mapping such that $\mu(R) = \mu(\phi_t(R))$ for all measurable $R \subseteq \Gamma$. The system's microstate at time t_0 (its 'initial condition'), $x(t_0)$, evolves into $x(t) = \phi_t(x(t_0))$ at time t. In a Hamiltonian system, energy is conserved, and hence the motion of the system is confined to the $6n - 1$–dimensional energy hypersurface Γ_E. The measure μ_L can be restricted to Γ_E, which induces a natural invariant measure μ on Γ_E.

To each macrostate M_i of the system, $i = 1, \ldots, m$, there corresponds a macroregion Γ_{M_i} consisting of all $x \in \Gamma_E$ for which the macroscopic variables assume the values characteristic for M_i. The Γ_{M_i} together form a partition of Γ_E, meaning that they do not overlap and jointly cover Γ_E up to measure zero. The Boltzmann entropy of a macrostate M_i is defined as $S_B(M_i) := k_B \log[\mu(\Gamma_{M_i})]$, where k_B is the Boltzmann constant. Given this, we define the Boltzmann entropy of a *system* at time t, $S_B(t)$, as 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 corresponding to $x(t)$.

Among the macrostates of a system, two are of particular importance: the equilibrium state, M_{eq}, and the system's state at the beginning of the process, M_p (also referred to as the 'past state'). The latter is, by assumption, a low-entropy state.[1] The idea now is that the behavior of $S_B(t)$ should mirror the behavior of the thermodynamic entropy S_{TD}, at least approximately.[2] So we expect the Boltzmann entropy of a system initially prepared in M_p to increase more or less monotonically, reach its maximum fairly quickly, and then remain at or near the maximum for a long time. In other words, we expect the dynamics to be such that it carries the system's initial state $x(t_0) \in \Gamma_{M_p}$ into $\Gamma_{M_{eq}}$ reasonably quickly and then keeps it there for a long time. I refer to this as 'thermodynamic-like behavior'.[3] The explanandum then is this: Why does the system under investigation behave in a thermodynamic-like way?

The standard Boltzmannian response to this question is to introduce probabilities and argue that the values of these probabilities come out such that the system is overwhelmingly likely to evolve in a thermodynamic-like way.[4] Typicality approaches to SM eschew commitment to probabilities and offer a different kind of explanation: the system behaves in a thermodynamic-like way because it is typical for systems of this kind to behave in this way.[5]

Before turning to a discussion of this approach, an important technical result needs to be stated. Under certain conditions, it is the case that $\Gamma_{M_{eq}}$ is vastly larger (with respect to μ) than any other macroregion. I refer to this matter of fact as the 'dominance of the equilibrium macrostate'. This dominance is then often glossed as being equivalent to the claim that for large n Γ_E is *almost entirely* taken up by equilibrium microstates (Bricmont 1996, 146; Goldstein 2001, 45; Zanghì 2005, 191, 196).

Some caution is needed here. In certain systems, nonequilibrium states can take up a substantial part of the phase space due to the degeneracy of below-equilibrium entropy values, and hence it is not true that Γ_E is almost entirely filled with equilibrium states, even if the equilibrium macrostate is by far the largest macrostate (Lavis 2005, 255–258; 2008, Section

1. If we study laboratory systems like the above-mentioned gas, M_p has low entropy by construction. If we take the universe as a whole to be the object of study, then that M_p is of low entropy is the subject matter of the so-called Past Hypothesis (Albert 2000, 96).

2. This 'mirroring' need not be perfect, and occasional deviations of the Boltzmann entropy from its thermodynamic counterpart are no cause for concern (Callender 2001).

3. This definition of thermodynamic-like behavior is the one adopted by those writing on typicality; see, e.g., Goldstein 2001, 43–44. Lavis (2005, 255) gives a somewhat different definition. These differences are inconsequential for what follows.

4. A discussion of the different ways of introducing probabilities into the Boltzmannian framework can be found in Frigg 2009.

5. This explanatory strategy is reminiscent of probabilistic explanations appealing to high probabilities and, hence, is open to similar objections. For the sake of argument, I set these worries aside and accept that something being typical has explanatory power.

2). However, it turns out that those nonequilibrium states that occupy most of the nonequilibrium area are close to equilibrium (in the sense of having close-to-equilibrium entropy values). We can then lump the equilibrium and these close-to-equilibrium states together and get an 'equilibrium or almost equilibrium' region, which indeed takes up most of Γ_E. The approach to equilibrium has then to be understood as the approach to this 'equilibrium or almost equilibrium' state, which is sufficient to give us thermodynamic-like behavior in the sense introduced above.

3. 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 a 'typicality measure'.[6] Roughly speaking, e is typical if most members of Σ have property P and e is one of them. More precisely, let Π be the subset of Σ consisting of all elements that have property P. Then the element e is typical if and only if (iff) $e \in \Pi$ and $\nu_\Sigma(\Pi) := \nu(\Pi)/\nu(\Sigma) \geq 1 - \varepsilon$, where $\varepsilon \geq 0$ is a small real number; $\nu_\Sigma(\cdot)$ is referred to as the 'measure conditional on Σ', or simply 'conditional measure'.[7] Derivatively, one can then refer to Π as the 'typical set' and to those elements that possess property P (i.e., the members of Π) as 'typical elements'. Conversely, an element e is atypical iff it belongs to the complement of Π, $\Omega := \Sigma/\Pi$, in which case we refer to Ω as the 'atypical set' and to its members as 'atypical elements'.

As an example, consider the number $\pi/4$, which is typical with respect to the interval [0, 1], the property 'not being specifiable by a finite number of digits', and the usual Lebesgue measure on the real numbers because it is a theorem of number theory that the set of all numbers that have this property has measure one. The element of interest in SM is a microstate x, and it is generally agreed that the relevant measure is the Lebesgue measure μ. However, views diverge when it comes to specifying the relevant set Σ and relevant property P.

I now turn to a discussion of three different typicality-based accounts of SM that emerge from the writings of Goldstein, Lebowitz, and Zanghì. In conversation, Goldstein and Zanghì have pointed out to me that they would not subscribe to Accounts 1 and 2 and that (something like) Account 3 is what they had intended. However, since the relevant papers can reasonably be read as proposing Accounts 1 and 2, it worth discussing them briefly to set the record straight (Sections 4 and 5) before turning to a detailed discussion of Account 3 (Section 6).

6. Typicality measures often are, but need not be, probability measures.

7. This definition of typicality is adapted from Dürr 1998, Section 2; Lavis 2005, 258; Zanghì 2005, 185; and Volchan 2007, 805.

4. First Account. The first account sets out to explain the approach to equilibrium in terms of the dominance of the equilibrium macrostate. Zanghì explains, "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; a fact often not understood by the critics of Boltzmann" (2005, 196; my translation).[8] On this view, then, a system approaches equilibrium *simply because* the overwhelming majority of states in Γ_E are equilibrium microstates. If we now associate Σ with Γ_E and property P with 'being an equilibrium state' (and, as indicated above, regard microstates as elements of interest and use the Lebesgue measure μ as a typicality measure), the dominance of the equilibrium macrostate implies that equilibrium microstates are typical and the view put forward in the above quote can be summarized as the claim that systems approach equilibrium because equilibrium microstates are typical and nonequilibrium microstates are atypical.

This explanation is unsuccessful. If a system is in an atypical microstate, it does not evolve into an equilibrium microstate *just because* the latter is typical. Typical states do not automatically attract trajectories.[9] In fact, there are Hamiltonians—for instance, the null Hamiltonian or a collection of uncoupled harmonic oscillators—that give rise to phase flows that do not carry nonequilibrium states into equilibrium. To explain why nonequilibrium microstates eventually wind up in equilibrium, the typicality of $\Gamma_{M_{eq}}$ is not enough, and appeal has to be made to the system's dynamics.

5. Second Account. A different line of argument can be found in Lebowitz 1993a, 1993b, 1999 and Goldstein and Lebowitz 2004. This account differs from the above in that it focuses on the internal structure of the microregions Γ_{M_i} rather than the entire phase space: "By 'typicality' we mean that for any [Γ_{M_i}] . . . the relative volume of the set of microstates [x] in [Γ_{M_i}] for which the second law is violated . . . goes to zero rapidly (exponentially) in the number of atoms and molecules in the system" (Goldstein and Lebowitz 2004, 57).[10]

This definition contains different elements that need to be distinguished. Let $\Gamma_{M_i}^{(++)}$ be the subset of Γ_{M_i} containing all x that lie on trajectories that come into Γ_{M_i} from a macrostate of higher entropy and that leave Γ_{M_i} entering into a macrostate of higher entropy; $\Gamma_{M_i}^{(+-)}$, $\Gamma_{M_i}^{(-+)}$, and $\Gamma_{M_i}^{(--)}$ are

8. Goldstein 2001, 43–44, 49, can be read as making a similar claim.

9. Uffink (2007, 979–980) illustrates this with the example of a trajectory.

10. Square brackets indicate that the original notation has been replaced by the notation used in this article. I use this convention throughout.

defined accordingly. These four subsets form a partition of Γ_{M_i}.[11] Furthermore, $\Gamma_{M_i}^{(+)} := \Gamma_{M_i}^{(++)} \cup \Gamma_{M_i}^{(-+)}$ and $\Gamma_{M_i}^{(-)} := \Gamma_{M_i}^{(+-)} \cup \Gamma_{M_i}^{(--)}$ are the subsets of Γ_{M_i} that have a higher- and lower-entropy future, respectively.

There is an interpretative question about how to understand the notion of a set of microstates in Γ_{M_i} violating the Second Law. A plausible reading takes these to be states that have an entropy-decreasing future. Let us call this property D. Hence, x has D iff $x \in \Gamma_{M_i}^{(-)}$. Entropy-decreasing states are atypical in Γ_{M_i} iff $\mu_i(\Gamma_{M_i}^{(-)}) < \varepsilon$, where $\mu_i(\cdot) := \mu(\cdot)/\mu(\Gamma_{M_i})$. Furthermore, let us say that a *system* has the property *global atypical entropy decrease*, GAD, iff entropy-decreasing microstates are atypical in every Γ_{M_i}. The claim made in the above quote then is tantamount to saying that a system with a reasonably large number of molecules is GAD.

This claim needs to be qualified. The atypicality of x with property D in Γ_{M_i} trivially implies $\mu(\Gamma_{M_i}^{(+-)}) < \varepsilon$. Due to the time reversal invariance of the Hamiltonian dynamics, we have $\mu(\Gamma_{M_i}^{(-+)}) = \mu(\Gamma_{M_i}^{(+-)})$, and therefore $\mu(\Gamma_{M_i}^{(-+)}) < \varepsilon$. Since $\Gamma_{M_i}^{(+)} = \Gamma_{M_i}^{(-+)} \cup \Gamma_{M_i}^{(++)}$, we obtain $\mu(\Gamma_{M_i}^{(++)}) > 1 - 2\varepsilon$. Hence, even if x with property D is atypical in Γ_{M_i}, it is not the case that, as we would expect, most states in Γ_{M_i} act thermodynamic-like since most states have a *higher* entropy *past*. But this is a familiar problem, and remedy can be found in conditionalizing on Γ_{M_p} (Albert 2000, Chapter 4).[12] So the correct requirement is that, at any time t, microstates violating the Second Law must be atypical in $\Gamma_{M_i} \cap \phi_t(\Gamma_{M_p})$ rather than only Γ_{M_i}.

Do relevant systems meet this requirement? Immediately after the passage quoted above, Goldstein and Lebowitz offer the following answer: "Boltzmann then argued that given this disparity in sizes of different M's, the time evolved $[M_{x(t)}]$ will be such that $[\mu(M_{x(t)})]$ and thus $[S_B(t)]$ will *typically* increase in accord with the law" (2004, 57). So the argument seems to be that the relevant condition must be true because the equilibrium state is much larger than other macrostates.

This is unconvincing. The disparity of sizes of macroregions is, of course, compatible with GAD, but the latter does not follow from the former. Whether macroregions have the above internal structure depends on the system's phase flow ϕ_t, and every attempt to answer this question without even mentioning the system's dynamics is doomed to failure right from the start (and this is true of both the qualified and the unqualified versions of the claim).

6. Third Account. As we have seen, the basic problem with the two ac-

11. For the sake of simplicity, I omit points lying on trajectories that stay at the same entropy level. Taking these into account would not alter the argument.

12. Notice that an attempt to define D in terms of $\Gamma_{M_i}^{(-)} \cup \Gamma_{M_i}^{(++)}$ rather than only $\Gamma_{M_i}^{(-)}$ leads to a contradiction.

counts discussed so far is that they attempt to explain the approach to equilibrium without reference to the system's dynamics. The third account, which emerges from a passage in Goldstein 2001, rectifies this problem: "[Γ_E] consists almost entirely of phase points in the equilibrium macrostate [$\Gamma_{M_{eq}}$], with ridiculously few exceptions whose totality has volume of order $10^{-10^{20}}$ relative to that of [Γ_E]. 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 [$\Gamma_{M_{eq}}$] and keeping it there for an extremely long time—unless, of course, [x] itself were ridiculously special" (Goldstein 2001, 43–44). This is an interesting claim but one that stands in need of clarification. A reasonable reading of this passage seems to be that an argument involving three different typicality claims is made:

> **Premise 1**. The system's macrostate structure is such that equilibrium states are typical in Γ_E in the sense introduced in Section 4.
> **Premise 2**. The system's Hamiltonian is typical in the class of all Hamiltonians.
> **Conclusion**. Initial conditions lying on trajectories showing thermodynamic-like behavior are typical in Γ_{M_p} with respect to $\mu_p(\cdot) := \mu(\cdot)/\mu(\Gamma_{M_p})$.

Let us refer to this as the 'T-Argument'. Premise 1 is familiar from Section 4 and is taken for granted here. Premise 2 and the conclusion are restatements in the language of typicality of the claims that the Hamiltonian of the system and the initial condition be not 'ridiculously special'.

The T-Argument, if sound, gives us the sought-after explanation of the approach to equilibrium in terms of typicality. But before we can address the question of soundness, we need to make Premise 2 more precise. What does it mean for a system's Hamiltonian to be typical in the class of all Hamiltonians? More specifically, what is the typicality measure, and what is the relevant property P?

Let us begin with the first question. The problem is that function spaces do not come equipped with normalized measures that can plausibly be used to capture the intuitive idea that some sets of functions are typical while others are atypical. A natural way around this difficulty is to replace the measure-theoretic notion of typicality introduced in Section 3 by a topological one based on Baire categories. Sets can be of two kinds: meager (first Baire category) or nonmeager (second Baire category). Loosely speaking, a meager set is the 'topological counterpart' of a set of measure zero in measure theory, and a nonmeager set is the 'counterpart' of a set of nonzero measure. Given this, it is natural to say that meager sets are atypical and nonmeager sets are typical. I call this notion of typicality 't-typicality' ("t" for "topological") and, to avoid confusion, from now on refer to the notion of typicality introduced in Section 3 as 'm-typicality', in order to make it explicit that it is a measure-theoretic notion.

Unfortunately there is no straightforward answer to the question about the property *P*. But a promising line of argument emerges from Dürr's (1998) and Maudlin's (2007) discussion of typicality in the so-called Galton Board, a triangular arrangement of nails on an infinitely long vertical board. Balls are fed into the board from the top and then move down the board. Every time a ball collides with a nail it moves either to the right (R) or to the left (L). If we follow a ball's trajectory and take down whether it moves to the left or to the right every time it hits a nail, we obtain a string of Rs and Ls that looks as random as one that has been generated by a coin toss: the Galton Board seems to exhibit random behavior. Why is this? Dürr's and Maudlin's answer is that the board appears random because random-looking trajectories are typical in the sense that the set of those initial conditions that give rise to nonrandom-looking trajectories has measure zero in the set of all possible initial conditions, and this is so because the board's dynamics is chaotic (Dürr 1998, Section 2).

Translating this idea into the context of SM suggests that the relevant property *P* is being chaotic. This sounds prima facie plausible and, most important, would make Premise 2 true. Markus and Meyer prove the following theorem:

> Completely integrable Hamiltonians are meager in the space of all normalized and infinitely differentiable Hamiltonians on a compact symplectic manyfold. (1974, 13)[13]

This implies that nonintegrable Hamiltonians are nonmeager, which is tantamount to saying that the class of chaotic Hamiltonians is nonmeager and, hence, t-typical.[14]

The question now is whether the T-Argument is valid, that is, whether the conclusion follows from the premises. This turns out to be a thorny issue. There is an entire class of systems that are chaotic but whose phase space is full of invariant curves, namely, so-called KAM systems (Argyris,

13. Two Hamiltonians that differ only by a constant are considered equivalent, and an equivalence class of Hamiltonians is called a 'normalized Hamiltonian'. This is because in practical calculations any Hamiltonian of this class can be chosen as a representative since they all yield the same flow (Markus and Meyer 1974, 11).

14. I here follow common practice and assume that nonintegrability implies chaos, at least on some region of the phase space. However, to the best of my knowledge there is no strict mathematical proof of this. Mathematicians refer to t-typical Hamiltonians as 'generic'.

Faust, and Haase 1994, Chapter 4). Naturally one would expect these curves to divide Γ_E into a set of closed volumes bounded by the invariant curves, which would prevent the system from approaching equilibrium (invariant curves are ubiquitous in KAM systems, and so it would be highly unlikely that Γ_{M_p} and $\Gamma_{M_{\mathrm{eq}}}$ would not be separated by one). In systems with 2 degrees of freedom, this is exactly what happens: the two-dimensional invariant surfaces divide the three-dimensional energy hypersurface into disconnected parts. Fortunately, the situation is better for systems with $f > 2$ degrees of freedom. The energy hypersurface has $2f - 1$ dimensions, and for another surface to divide it into two disconnected parts, it must have $2f - 2$ dimensions. But the invariant KAM tori are f-dimensional, and since $2f - 2 > f$ for all $f > 2$, invariant KAM surfaces do not divide Γ_E into separate parts; the invariant surfaces are a bit like lines in a three-dimensional Euclidean space. So the trajectories can, in principle, wander around relatively unhindered and without being 'sandwiched' between invariant surfaces. This process is known as Arnold Diffusion. It has first been proven analytically to exist in a particular example, and there is now numerical evidence that it can also be found in other systems. In such systems, the chaotic parts of Γ_E are connected and form a single net, the so-called Arnold Web, which permeates the entire phase space in the sense that a trajectory moving on the web will eventually visit almost every finite region of Γ_E.[15]

This looks like what we need, but unfortunately some difficulties arise on the finishing line. These can be circumvented only at the cost of accepting three conjectures, which are supported only by plausibility arguments. First, there is no proof for the existence of Arnold Webs in *all* nonintegrable systems with $f > 2$. The good news is that so far there are no known examples in which this is not the case, and so we can venture the conjecture that all nonintegrable systems with $f > 2$ have Arnold Webs (Conjecture 1).[16]

Second, there is a question about the relative measure in Γ_E occupied by Arnold Webs. It could in principle be that these webs are of measure zero or else fill only a small part of the phase space. If this were the case, it would be unlikely that m-typical initial conditions would come to lie on trajectories that wander around randomly (and therefore wind up in $\Gamma_{M_{\mathrm{eq}}}$), which would undercut the conclusion in the T-Argument. However, numerical simulations on simple systems have shown that the relative measure occupied by invariant KAM curves decreases as f increases (Ear-

15. In fact, Lichtenberg and Liebermann (1992, 61) and Ott (1993, 257) say that the system visits every finite region of Γ_E, but this seems to be too strong.

16. Or if not *all* nonintegrable systems have Arnold Webs, then those that do not should be so few that the class of those with Arnold Webs is still nonmeager.

man and Rédei 1996, 70). Furthermore, Sklar (1993, 175) observes that there are good numerical reasons to think that large systems are "at least ergodic-like" on the "overwhelmingly largest part" of the accessible phase space, and Vranas (1998, 695–698) gathers a welter of numerical evidence for the conclusion that many systems of interest in SM are ε-ergodic, that is, ergodic on nearly the entire energy hypersurface. This suggests that it may well be the case not only that Arnold Webs have finite measure but that they in fact fill most of Γ_E (Conjecture 2).

Third, in order to explain thermodynamic-like behavior, we need to know how much time the system spends in different parts of the phase space. Again, little is proven rigorously, but Ott (1993, 257) suggests that the system is ergodic on the Arnold Web (Conjecture 3). The numerical evidence just mentioned supports this conjecture.

If these three assumptions are correct, then the T-Argument is sound. By Premise 2 the system is chaotic, and by Conjecture 1 it has an Arnold Web, which, by Conjecture 2, fills most of the energy surface and hence most of Γ_{M_p}. Therefore, points on the web are m-typical in Γ_{M_p}. By Conjecture 3, these points wander around ergodically on the web and hence approach $\Gamma_{M_{\text{eq}}}$ fairly soon and stay there for a long time (where "fairly soon" means that the time taken to arrive at equilibrium is much shorter than the time spent in equilibrium) because, by Premise 1, nonequilibrium states occupy a much smaller volume than equilibrium states.

To put this argument on secure footing, more would have to be said about the three conjectures. This is an extremely difficult task, and so it is worth asking whether there is not a simpler way to arrive at the same conclusion. I now discuss a plausible way of doing so and show that it is a blind alley. Hence, there is no way around trying to make progress on the conjectures.

The new line of argument departs from the observation that we might have chosen too liberal a notion of chaos. In fact, there is a great deal of controversy over the correct characterization of chaos (Smith 1998, Chapter 10), and so we might say that KAM systems show the wrong kind of chaotic behavior: they exhibit 'local chaos', meaning that the dynamics are chaotic only on parts of the phase space. What we need, so the argument goes, is that the relevant systems show 'global chaos', which disqualifies KAM systems.

The question then becomes how to characterize global chaos. Commonly this is done in either of two ways: a topological way and a measure-theoretic one. The latter always involves ergodicity and is therefore untenable: Markus and Meyer (1974, 14) also prove that in the space of all normalized and infinitely differentiable Hamiltonians on a compact symplectic manyfold, the class of ergodic Hamiltonians is meager, and hence strongly chaotic systems are t-atypical. Topological definitions of chaos (the best known of which is Devaney's) always involve topological transitivity, the condition that for any two open regions A and B in Γ_E, there is a trajectory initiating in A that eventually visits B. But this condition does not fit the bill: while it is at least plausible that topological chaos is a sufficient condition for the approach to equilibrium, it requires a revision of P that seems to render Premise 2 false. As Markus and Meyer (1974, 1) point out, ergodic systems are meager because generic systems have invariant surfaces preventing the trajectory from accessing the entire phase-energy hypersurface. But a system that cannot access certain regions of Γ_E not only fails to be ergodic; it also fails to be topologically transitive. So topologically transitive systems must be meager too and, hence, also fail to be t-typical. For this reason, shifting attention to global chaos is a dead end.

7. Conclusion. I have distinguished three different accounts of how typicality is used to explain thermodynamic-like behavior. I have argued that while the first two fail, the third is promising, and I have sketched a proof. The proof rests on three conjectures that need to be further substantiated to put the argument on secure footing. Furthermore, an argument needs to be given that m-typicality and t-typicality have explanatory power from the point of view of physics.

REFERENCES

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

Argyris, John, Gunter Faust, and Maria Haase (1994), *An Exploration of Chaos*. Amsterdam: Elsevier.

Bricmont, Jean (1996), "Science of Chaos or Chaos in Science?", in Paul R. Gross, Norman Levitt, and Martin W. Lewis (eds.), *The Flight from Science and Reason*. Annals of the New York Academy of Sciences 775. New York: New York Academy of Sciences, 131–175.

Callender, Craig (2001), "Taking Thermodynamics Too Seriously", *Studies in History and Philosophy of Modern Physics* 32: 539–553.

Dürr, Detlef (1998), "Über den Zufall in der Physik", paper given at the 1998 Leopoldina Meeting, Halle, http://www.mathematik.uni-muenchen.de/%7Eduerr/Zufall/zufall.html.

Earman, John, and Miklós Rédei (1996), "Why Ergodic Theory Does Not Explain the Success of Equilibrium Statistical Mechanics", *British Journal for the Philosophy of Science* 47: 63–78.

Frigg, Roman (2009), "Probability in Boltzmannian Statistical Mechanics", in Gerhard Ernst and Andreas Hüttemann (eds.), *Time, Chance and Reduction: Philosophical Aspects of Statistical Mechanics*. Cambridge: Cambridge University Press.

Goldstein, Sheldon (2001), "Boltzmann's Approach to Statistical Mechanics", in Jean Bricmont, Detlef Dürr, Maria Carla Galavotti, Gian Carlo Ghirardi, Francesco Petruccione, and Nino Zanghì (eds.), *Chance in Physics: Foundations and Perspectives*. Berlin: Springer, 39–54.

Goldstein, Sheldon, and Joel L. Lebowitz (2004), "On the (Boltzmann) Entropy of Non-equilibrium Systems", *Physica D* 193: 53–66.

Lavis, David (2005), "Boltzmann and Gibbs: An Attempted Reconciliation", *Studies in History and Philosophy of Modern Physics* 36: 245–273.

——— (2008), "Boltzmann, Gibbs, and the Concept of Equilibrium", *Philosophy of Science* 75: 682–696.

Lebowitz, Joel L. (1993a), "Boltzmann's Entropy and Time's Arrow", *Physics Today* (September): 32–38.

——— (1993b), "Macroscopic Laws, Microscopic Dynamics, Time's Arrow and Boltzmann's Entropy", *Physica A* 194: 1–27.

——— (1999), "Statistical Mechanics: A Selective Review of Two Central Issues", *Reviews of Modern Physics* 71: 346–357.

Lichtenberg, Allan J., and M. A. Liebermann (1992), *Regular and Chaotic Dynamics*. 2nd ed. Berlin: Springer.

Markus, Lawrence, and Kenneth R. Meyer (1974), "Generic Hamiltonian Dynamical Systems Are neither Integrable nor Ergodic", *Memoirs of the American Mathematical Society* 144: 1–52.

Maudlin, Tim (2007), "What Could Be Objective about Probabilities?", *Studies in History and Philosophy of Modern Physics* 38: 275–291.

Ott, Edward (1993), *Chaos in Dynamical Systems*. Cambridge: Cambridge University Press.

Sklar, Lawrence (1993), *Physics and Chance: Philosophical Issues in the Foundations of Statistical Mechanics*. Cambridge: Cambridge University Press.

Smith, Peter (1998), *Explaining Chaos*. Cambridge: Cambridge University Press.

Uffink, Jos (2007), "Compendium of the Foundations of Classical Statistical Physics", in Jeremy Butterfield and John Earman (eds.), *Philosophy of Physics*. Amsterdam: North-Holland, 923–1047.

Volchan, Sérgio B. (2007), "Probability as Typicality", *Studies in History and Philosophy of Modern Physics* 38: 801–814.

Vranas, Peter B. M. (1998), "Epsilon-Ergodicity and the Success of Equilibrium Statistical Mechanics", *Philosophy of Science* 65: 688–708.

Zanghì, Nino (2005), "I fondamenti concettuali dell'approccio statistico in fisica", in Valia Allori, Mauro Dorato, Federico Laudisa, and Nino Zanghì (eds.), *La natura delle cose: Introduzione ai fundamenti e alla filosofia della fisica*. Rome: Carocci.