How Typical!

An Epistemological Approach to Typicality in Statistical Mechanics

Massimiliano Badino

Universitat Autònoma de Barcelona – Massachussets Institute of Technology

DRAFT

Abstract

The recent use of typicality in statistical mechanics for foundational purposes has stirred an important debate involving both philosophers and physicists. While this debate customarily focuses on technical issues, in this paper I try to approach the problem from an epistemological angle. The discussion is driven by two questions: (1) What does typicality add to the concept of measure? (2) What kind of explanation, if any, does typicality yield? By distinguishing the notions of 'typicality-as-vast-majority' and 'typicality-as-best-exemplar', I argue that the former goes beyond the concept of measure. Furthermore, I also argue that typicality aims at providing us with a form of causal explanation of equilibrium.

1 Introduction

Among the recent foundational approaches to equilibrium statistical mechanics, typicality has stirred the most heated and interesting debate. While fiercely supported by some prominent physicists (Dürr, Goldstein and Zanghì 1992, Lebowitz 1993b, Lebowitz 1999, Dürr 2001, Goldstein 2001, Zanghì 2005, Goldstein 2012), it has been repeatedly challenged by philosophers of science (Frigg 2007, Frigg 2009, Hemmo and Shenker

2012, Pitowsky 2012). Roughly speaking, the typicality approach aims at explaining the qualitative features of thermal equilibrium—its being unidirectional and exceptionless— by showing that these features are 'typical' of usual statistico-mechanical systems. The notion of typicality is notoriously mudded. It is related to probability and measure, but, in the intention of the 'typicalists', to say that a property or a behavior is typical means something more than 'highly probable'.

Criticisms against this approach have focused especially on the physical, mathematical, and philosophical acceptability of its steps. Beneath its apparent intuitiveness, typicality reveals problematic assumptions. In the attempt to unravel the technical intricacies nested in the notion of typicality, philosophers have mostly—and wittingly— left aside the epistemological issue of the explanatory value of this approach. Itamar Pitowsky was possibly the first to adumbrate this issue when compared typicality and Lebesgue measure and concluded that the latter seems to have all the virtues of the former, without having its problems (Pitowsky 2012, 54-56). This raises the question: What is the explanatory surplus attached to typicality? This question goes hand in hand with a more general one: What kind of explanation is incapsulated in the use of typicality made by philosophers and physicists? These questions help us see the debate on typicality from a new, epistemological perspective.

My strategy to tackle them consists of three steps. First, I discuss the structure of the typicality explanation used in statistical mechanics. This explanation relies on two typicality-claims concerning the phase space and the Hamiltonian of the system. The explanation is successful if one can characterize the latter in order to use the former to conclude the qualitative features of equilibrium. Second, I analyze the semantics of typicality-claims and I argue that they can be dually spelled out in terms of properties possessed by many individuals (typicality-as-vast-majority) or in terms of the fitting with the process generating the kind of individuals having the relevant properties

2

(typicality-as-best-exemplar). Third, I show that current positions on typicality can be categorized according to their reading of the typicality-claims. My conclusions are that (1) typicality does not differ from measure insofar as it is interpreted as typicality-as-vast-majority and (2) typicality yields a kind of causal explanation of equilibrium.

2 Epistemic Stories and Typicality-Claims

My epistemological analysis relies on two resources. First, for reasons that will become clear later, it is useful to think of typicality within the framework of a manipulative view of explanation. Thus, I assume that to explain something, in this case our experience of equilibrium, means to provide information on difference-making factors that enables us to potentially control the occurrence of the explanandum. To put it in other terms, to explain means to answer a what-if-the-things-had-been-different type of question.[1] The answer to such a question is an epistemic story constituted by variables, assumptions, and invariant relations arranged in a certain argumentative pattern. Some of these variables and assumptions may serve the auxiliary purpose of setting the context, while others are difference-making elements in the production of the explanandum. These variables are the causal factors and their relation to the explanandum is discussed in detail later on.

My second resource to evaluate the explanatory status of typicality is semantic analysis. There is a considerable amount of uncertainty about what logical or linguistic objects typicality refers to. A quick glimpse at the literature shows that it is generally interpreted as a feature of a property or a behavior. In the attempt to capture both uses at the same time, I will focus on sentences in which typicality is ascribed, i.e.,

---

[1]Here, I refer especially to James Woodward's thorough treatment in (Woodward 2000, Woodward 2003).

typicality-claims. Typicality-claims have the following forms:

(1) Typically, elements x of the set K have (the property or behavior) P.

A typicality-claim (1) is equivalent to saying that the property or behavior is typical in K, but is easier to analyze. We will see below that discussing typicality-claim allows us to draw interesting conclusions from the semantics of these particular statements.

Before proceeding, a word of caution is in order. In this paper, I am concerned first and foremost with the debate on typicality in statistical mechanics. I expect some aspects of my analysis to have a bearing on typicality explanation at large—if there is such a species. This is the case, I think, for the semantics of typicality-claims. However, in the following sections, I only treat the form that typicality assumes in statistical mechanics and the conclusions on its explanatory value should not be generalized too quickly to other cases.

3   The Structure of the Typicality Explanation

The first legitimate question to ask is: What typicality is actually supposed to explain? A clarification of this aspects can be found in Joel Lebowitz's lucid article (Lebowitz 1993b). In the opening section of his paper, Lebowitz distinguishes neatly between qualitative and quantitative aspects of equilibrium. Firstly, he introduces the concepts of macro- and microstate. A microstate is a complete description of the state of the system in terms of the positions and momenta of each particle. If the system consists of N particles, its microstate x is a point in a $6N$-dimensional phase space. Microstates can be grouped into macrostates, i.e., regions in the phase space comprising microstates that look the same from a macroscopic point of view.

Ludwig Boltzmann showed that one can attach a quantity, called Boltzmann entropy

S_B, to a macrostate M and consequently to each of its microstates. Lebowitz's claim is that "S_B typically increases in a way which explains and describes qualitatively the evolution towards equilibrium of such systems" (Lebowitz 1993b, 2). What the typical behavior of entropy is supposed to explain is precisely our day-to-day experience of equilibrium. He clarifies further this point in the following section. Let us assume to have a container divided into two chambers by a partition and assume that all particles are contained in one of the two chambers. Let us remove the partition so that the particles move freely from one chamber to the other until they fill up the container uniformly. Let also assume that we have taken a series of snapshots of this process, from the initial state to the final, uniform distribution. If asked to order the snapshots, "the 'obvious order', based on experience" is that the system moves from a non-uniform to a uniform state. This is what our experience of equilibrium amounts to and this is what typicality is supposed to explain. There are also quantitative aspects of equilibrium, but they require different sorts of resources: "the quantitative description of the macroscopic evolution is given by hydrodynamical-type equations which can be derived (explicitly, in some cases) from the microscopic dynamics by utilizing the collective aspect of macrobehavior, i.e., as a law of large numbers arising from the very large macro/micro-ratio." The quantitative aspects are related to typicality, but are not directly explained by it. Hence, typicality explains the thermodynamic-like behavior, i.e., the time increase of the Boltzmann entropy: $S_B(t) \leq S_B(t+1)$.

Let us now analyze the epistemic story of typicality. First, the story comprises a set of concepts and assumptions of standard statistical mechanics. Let Σ be the accessible phase space of a physical system S, x_i the phase points of this space corresponding to the possible microstates of S. Let us assume that the space can be partitioned into a series of disjoint sets $M_1, M_2, \ldots, M_E$. To each microstate $x_i \in M_i$ one can assign a Boltzmann entropy $S_B(M_i)$. The Boltzmann entropy $S_B(M_E)$ for equilibrium is a maximum. Let the dynamics be given by an Hamiltonian flow $H(x)$, i.e., $H(x_{t-1}) = x_t$. A

trajectory of the system is a sequence of microstates x_i of S generated by $H(x)$. These ingredients are then arranged in the following argumentative pattern of typicality:[2]

(T1) Typically, a microstate x_i belongs to M_E, therefore, typically, $S_B(x_i) = S_B(M_E)$.

(T2) Typically, the Hamiltonian flow $H(x_t)$ originates a trajectory that has a certain property Γ.

(C) The property Γ together with (T1) entails that S exhibits a thermodynamic-like behavior.

The claim (T1) expresses a well-defined and undisputed fact of the phase space, i.e., that an enormous amount of microstates belong to the equilibrium region. Although well-founded, (T1) concerns the structure of the phase space, but it does not tell us anything specific about the behavior of the physical trajectory of the system over time, which is determined by the initial conditions and the Hamiltonian. The typicality-claim (T2) deals precisely with these elements. The crux of the argument is the characterization of the property Γ, which leads to the thermodynamic-like behavior on the basis of what asserted in (T1). It is because of Γ that the overall structure of the phase space—in itself a geometrical fact—matters for the behavior of physical systems in time. Thus, the essence of the typicality approach and its explanatory value depend precisely on this mysterious property Γ. A semantic analysis of typicality-claims can help us understand the attempts done by philosophers and physicists to characterize Γ and, consequently, to evaluate the explanatory value of typicality. First of all, I argue that, in general, there are two dual ways to spell out a typicality-claim. Next, I discuss which reading has been associated to (T1) and (T2). We will see that, while in statistical mechanics there is a unanimous consensus about how (T1) should be spelled out, the status of

---

[2]This argumentative pattern tries to capture the variegated uses of typicality in physical and philosophical literature. For some writers, the essence of the argument consists in (T1) only, but, as Roman Frigg has argued, some dynamical step of the sort of (T2) is also required (Frigg 2009, Frigg and Werndl 2012).

(T2) is much less clear.

4 Typicality-1 and Typicality-2

A typicality-claim is a proposition that would elicit a reaction such as: "well, how typical!". Here I am concerned with the kind of situations in which we utter the how-typical reaction. One situation in which this can happen is when we deal with a property or behavior represented by the vast majority of the individuals under examination. I buy a ticket of the National Lottery and it turns out that I don't win any prize: how typical. This situation is described by the following statement:

(2)   Typically, a ticket of the National Lottery is a non-winning.

One obvious way to interpret (2) is that the vast majority of lottery tickets do not win any prize, because there are very many tickets and very few prizes. In this case, (2) is taken to mean typicality-as-vast-majority or 'typicality-1' for short. Typicality-1 commits us to define:

(C1)   a set of individuals K (lottery tickets),

(C2)   a property P (being a non-winning ticket),

(C3)   a procedure μ to count how many individuals of K have P.

(C1)-(C3), in turn, allows us to define the truth-condition of the typicality-claim (2) interpreted as typicality-1:

(TC1)   A typicality-claim is true if and only if the vast majority of the members of K have P when counted according to the procedure μ.

Let me highlight two points about (2) and its truth-condition. First, (TC1) is vague because there is no threshold separating a simple majority from a vast majority. As it is often the case with vague concepts, we may have clear examples and counterexamples of vast majority, but we are unable to fix a discriminating value. Thus, typicality-1 always involves a certain amount of conventionality. Second, (TC1) depends on a counting procedure and not on probability, at least directly. In other words, what makes (2) true is not the low probability of attaining a winning ticket, but rather the fact that, according to a certain counting procedure, there are very many non-winning ticket. Granted, in some circumstances there is a straightforward relation between counting and probability or, better, between measure function and probability function. However the two concepts must be kept distinct for a number of reasons. For one, there are several definitions of probability and it is unclear which one should be attached to a certain counting procedure. More importantly, probability involves assumptions on the process. For instance, while the counting of black balls in a box is a fairly unproblematic procedure, the calculation of the probability of extracting a black ball, albeit related with their number, calls for assumptions on the drawing process. I will return to this important difference and its consequences for statistical mechanics in Section 5.1.

Typicality-1 is not the only way to read a typicality-claim. To appreciate this point, let us look again at the conditions (C1)-(C3). The truth-condition (TC1) demands that we count the members of the set K by means of the procedure μ and establishes that the vast majority have P. But to count the members of a certain set, one needs an independent way to define the set. For us to ascertain that the vast majority of elements x of K have P, we need to know what is for an individual x to be an element of K. I call the conditions that make an individual x a member of K the "process" producing x's of K. Thus, one can replace (C3) above with:

(C4)   A process Π to define members of K.

For instance, for an object to be a lottery ticket, a series of conditions must be realized: it must be printed by a certain agency, it must have a number assigned, there must be a drawing and prizing procedure and so on. Now (C1), (C2), and (C4) lead to the following truth-condition:

(TC2) A typicality-claim is true if and only if the best exemplars of K have the property P.

Let me clarify (TC2) with an example. I meet a professional piano player and I note that her hands are quick and agile: how typical. In fact:

(3) Typically, a professional piano player has agile hands.

Claim (3) can be spelled out in terms of typicality-as-vast-majority: we can ideally check the hands of all professional piano players and find out the most of them have agile hands. However, we can also read (3) as meaning that the specific training piano players go through enhances the agility of their hands. It is a typical feature of their profession. This alternative reading suggests that the typicality-claim refers to the best exemplar of a process through which an individual becomes member of a certain set, in this case, the set of professional piano players. The typical member of the set is an individual that embodies certain features that "best fit" the process; and if it does, it is the best exemplar of that set. I call this way of interpreting a typicality-claim 'typicality-as-best-exemplar' or 'typicality-as-best-fit' or typicality-2 for short. The vagueness of "vast majority" is here replaced by the also vague concept of "best exemplar". Consider the process of flipping a fair coin 10,000 times. A "best exemplar" sequence is not a sequence with precisely 5,000 heads, but any sequence in which heads and tails show up

roughly half of the time.[3] Again, there is no clear-cut discrimination value to define what "roughly" means. The determination of the conditions for a sequence or an individual to be a best exemplar is an open question. Some attempts in this direction have been carried out in the discussion of a Humean approach to lawhood and counterfactuals. Let me briefly summarize this analysis.

From the perspective of Humean supervenience, a law of nature is a regularity that fits the spatio-temporal facts of the world. One of the most popular approaches within this framework is David Lewis' best system analysis (Lewis 1994). According to Lewis, contrary to an accidental generalization, a law of nature is part of a deductive system, which describes the observed regularities and strikes the best balance between simplicity, strength, and fit. Of these three virtues, the third one is the most difficult to formulate. Intuitively, the best system is the one assigning the highest probability to the actual course of observed events. But if the world is intrinsically chancy, this proposal runs into the so-called zero-fit problem. Let us assume a world in which the one and only regularity is a sequence of heads and tails. Ideally, our criterion should lead us to select as best system the one assigning a probability distribution equal to the observed ratio of heads and tails. But, independently of the system, the number of possible sequences grows as 2^n and the probability given by any system to the only actual sequence decreases accordingly. As a consequence, any system would give zero-probability to the observed facts of the world.

Chancy worlds cause troubles also to Lewis' semantics of counterfactuals. It is well known that, for Lewis, a counterfactual statement such as "if it were the case that F, it would be the case that G" is true if and only if it is the case that G in all possible F-worlds closest to the actual world (Lewis 1973). Among the conditions that define

---

[3]We also have to add other properties. A sequence in which the first 5,000 flips come out head and the second 5,000 tail is obviously not typical. More on this point below.

the similarity between possible worlds, the most important are the respect of the laws of nature and the maximization of the spatio-temporal regions in which there is a perfect match with experience. But if the sequence of events is intrinsically probabilistic, one can satisfy both conditions and still create embarrassing counterfactuals. In particular, we can select a world that is completely identical to ours except for a localized event and construct therefore a true counterfactual. For instance, let assume that in this world I spend a whole week in Rio de Janeiro and, on Wednesday, I buy a parrot. We can imagine a world identical to the actual one in which, on Tuesday, instead of roaming around in Rio, I drop a plate in Beijing. In a chancy world, the counterfactual "If I dropped a plate in Beijing on Tuesday, I would buy a parrot in Rio on Wednesday" comes out true. The gist of the trouble is that, according to this procedure, a counterfactual is true if the consequent is true in the actual world, independently of the antecedent.

The two problems are obviously related. Lewis' conception of lawhood cannot come to the rescue of counterfactuals because, as we have seen, in a chancy world any best system gives zero probability to any sequence, even the one in which I drop a plate in Beijing on Tuesday and buy a parrot in Rio on Wednesday. Another way to see the connection is the following. The reason of the foregoing problem with counterfactuals is that, in a chancy world, even the best system cannot rule out highly improbable segments of events. For instance, even if it is true that the probability of head is $1/2$, it remains possible—albeit hugely improbable—to observe a sequence of 10,000 consecutive tails. But this is also what the zero-fit problem boils down to: the possible systems are compatible with so many sequences that eventually they will all assign zero probability to each.

Recent attempts at solving these two interconnected problems replace the inadequate notion of probabilistic fit with typicality. Within the framework of best system analysis,

Adam Elga has adopted a concept of typicality derived from the work of Haim Gaifman and Marc Snir (Gaifman and Snir 1980, Elga 2004). Details are complicated, but the basic idea is very familiar to practitioners of statistical mechanics. Elga suggests that the fitting of a deductive system with the observed regularities should not be measured in terms of probability yielded by the system to the regularity, but in terms of certain properties instantiated by the regularity and which have high probability according to the system. Put in other terms, given a deductive system, the typical regularity according to that system is the regularity that instantiates a series of properties that, if one adopts that system, have high probability. For instance, if the deductive system ascribes probability 1/2 to head, a high-probability property is "to have the same number of heads and tails" and a sequence is typical if it instantiates this property. According to Elga, if we define the best system as the system for which the observed regularities are typical, we do not incur in the zero-fit problem (Elga 2004, 71-72). This idea has been applied to the semantics of counterfactuals by J. Robert Williams (Williams 2008, Dodd 2011). William's proposal consists in adding to Lewis' conditions of similarity between possible words the requirement that close worlds do not present even localized atypical events, i.e., events that do not fit the existing laws of nature according to a certain set of high-probability properties.

The concept of typicality-2 takes the inverse path as typicality-1. In the case of typicality-1, one starts with a clear definition of what is the typical property that a behavior should instantiate. In statistical mechanics, this is the property of belonging to the equilibrium phase region. According to typicality-2, instead, one begins with an actual typical behavior, which is such relatively to "test properties" that must be instantiated by it. The open question is how to characterize the test properties. Elga, for instance, suggests that one should use simple properties formulated by means of a selected number of predicates in a particularly simple language. This requirement is meant to rule out the trivial 'exact distribution' property which has the highest probability for the actual

world, but is obviously very complicate to describe.

These proposals aim at making logical and epistemological sense of the concept of typicality-as-best-fit. This concept tries to capture the intuition that, if a behavior is typical, then it is a feature of the best fitting exemplar of the underlying process. Accordingly, a professional piano player with agile hands is the perfect exemplar of piano training process. It is important to realize that typicality-1 and typicality-2 are not simply alternative ways to read typicality-claims, but they are dual. Typicality-1 is typicality of properties, while typicality-2 is typicality of behavior. Typical properties are instantiated by a certain behavior and a typical behavior is such respect to certain properties. Ideally, (C1)-(C4) are all relevant to a complete claim of typicality. As a matter of fact, however, one reads a typicality-claim either as typicality-1 or as typicality-2. To chose one of the other reading depends largely on the conditions of the problem. Typicality-1 is certainly easier to realize, but it is also less informative from a causal and explanatory standpoint. Take again the typicality-claim (3). It is easy to count over the set of piano players, but this counting would give us only a sort of statistical correlation. By contrast, typicality-2 is more difficult to realize, but it furnishes more substantial explanatory information. It might be more complicated to relate the piano training process and the property of having agile hands, but if we manage to do it convincingly, the result would be much more rewarding in terms of explanation.

These considerations are very general and based on the semantics of typicality-claims. The actual circumstances in which these claims are held can be further complicated and make the evaluation of the explanatory information more nuanced, as we will see in the case of statistical mechanics. But the crucial point concerns the duality. Although different, these two interpretations are intertwined. In many circumstances, the two aspects are fairly separated, because one only needs minimal information to define K. In the case of (3), for instance, K is easily defined as the set of individuals who play pi-

ano as their profession. This establishes unambiguously the set over which we have to count. But consider the following example:

(4) A typical athlete is a well-trained person.

This typicality-claim is more problematic than (3). In terms of typicality-1, (4) means that if we examine the set of athletes and count over them, it turns out that the vast majority of the members of the set is well-trained. But how are we to define the set of athletes? Athletes are persons who conduct a certain life, who undergo a certain process whose result is, typically, to produce a well-trained body. Thus, if (4) is taken to mean that the vast majority of athletes are well-trained (typicality-1), this immediately refers to the definition of an athlete as a person, who undergoes a certain process, which produces a well-trained physique as best exemplar (typicality-2); while if (4) is taken to mean that the typical athlete fits a style of life whose result is to produce a well-trained physique (typicality-2), this immediately leads to the conclusion that the vast majority of the individuals following that style of life have a well-trained physique (typicality-1). Typicality-1 and typicality-2 are inextricably blended. In the remainder of the paper, I explore the consequences of this duality for statistical mechanics.

5 The Explanatory Value of Typicality

5.1 (T1) as typicality-1

There is a general consensus that (T1) should be spelled out as typicality-1. In fact, typicality-1 seems to be the favorite interpretation for typicality-claims among physicists and philosophers of science. Most of the definitions of the concept of typicality make reference to a counting procedure and to the resulting overwhelming majority of

typical results. For instance, Roman Frigg defines typicality in the following way:

> Intuitively, something 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. (Frigg 2009, 997-998)

Although this quote mixes typicality-claims that can be legitimately spelled out differently, there can be no doubt that philosophers and physicists *typically* read typicality-claims in this way. Another notable example, even more straightforwardly related to the case of statistical mechanics, is the following:

> Generally speaking, a set is typical if it contains an 'overwhelming majority' of points in some specified sense. In classical statistical mechanics there is a 'natural' sense: namely sets of full phase-space volume. (Volchan 2007, 803)

Lebowitz also endorses this reading:

> Typical, as used here, means that the set of microstates corresponding to a given macrostate M for which the evolution leads to a macroscopic decrease in the Boltzmann entropy during some fixed time period τ, occupies a subset of [the accessible phase space] whose Liouville volume is a fraction of [the volume of the phase space] which goes very rapidly (exponentially) to zero as the number of atoms in the system increases. (Lebowitz 1999, S348)

In this quote, Lebowitz also mentions the role of the number of degrees of freedom on which I will come back soon. Thus, there is a virtually unanimous agreement that (T1) should be taken to mean that the vast majority of microstates x_i belong to the macrostate M_E. Two main consequences have followed from this agreement. First, most of the criticisms against typicality have concentrated on the concept of measure. One matter of concern, for instance, is that although measure and probability are closely connected, they still convey different kinds of information. Measure deals with the structure of the ideal phase space, while probability has to do with the relative frequency of

physical properties. The leap from one to the other can be problematic. For instance, it is well known that zero-measure does not necessarily mean zero-probability, i.e., impossibility, because even an infinite set can have zero-measure according to a suitable measure. A related difficulty is the justification of the measure function. In statistical mechanics, it is customary to use the Lebesgue measure, but what argument can be offered for this choice? Meir Hemmo and Orly Shenker have argued that it is illegitimate to conclude something about probability—solidly rooted into observed relative frequencies—on the basis of the a priori choice of the Lebesgue measure and have suggested that the "epistemological arrow" should be inverted: one should determine the measure function on the basis of the observed physical probability (Hemmo and Shenker 2012, 182-191).

The second consequence concerns the explanatory information conveyed by (T1) read as typicality-1. Clearly, the message of the typicality-claim (T1) is that equilibrium happens because the equilibrium macrostate is dominant respect to other states. Typicality is not just a matter of how big the phase space volume is, but rather how 'much bigger' is respect to the alternative: "insofar as typicality is concerned [. . .] all that matters is which sets have very large measure and which very small" (Goldstein 2012, 66). It must be noted that the dominance of the equilibrium state depends crucially on a physical feature of the systems, i.e., the high number of degrees of freedom. It is precisely the astronomically large number of degrees of freedom that marks the difference in size and therefore the huge extension of the equilibrium region. Lebowitz explicitly stresses this point when he states:

> [T]he central role in time asymmetric behavior is played by the very large number of degrees of freedom involved in the evolution of macroscopic systems. It is only this which permits statistical predictions to become "certain" ones for typical individual realizations, where, after all, we actually observe irreversible behavior. Thus typicality is very robust—the essential features of macroscopic behavior are not dependent on any precise assumptions, such as ergodicity, mixing or "equal a priori probabilities", being

strictly satisfied by the statistical distributions. (Lebowitz 1993b, 3)

From this quote it is clear that the high number of degrees of freedom affects the occurrence of equilibrium, but how are we to express this in explanatory terms? I submit that, if we adopt a manipulative view of causality, we can conceptualize the number of degrees of freedom as playing the role of causal variable. Woodward's account of "C causes E" demands that: (i) C and E can be represented by variables; (ii) only an intervention on the value of the variable C results in a certain change in the value of the variable E, and (iii) the relation between variables C and E remains the same, i.e., it is invariant. Briefly said, "an explanation ought to be such that it can be used to answer what I call a *what-if-things-had-been-different* question" (Woodward 2003, 11). Let us discuss how the number of degrees of freedom fits Woodward's requisites.

First of all, the number of degrees of freedom acts as causal factor for equilibrium because the increase or decrease of that number makes the ratio of the equilibrium volume to the other volumes larger or smaller. Both the number of degrees of freedom and the equilibrium can therefore be quantified and variables can be attached to both. Secondly, Woodward defines an intervention on some variable X with respect to some second variable Y as "a causal process that changes X in an appropriately exogenous way, so that if a change in Y occurs, it occurs only in virtue of the change in X and not as a result of some other set of causal factors" (Woodward 2000, 199-200). An intervention does not need to be actually carried out by a human agent. It suffices to specify, possibly in a counterfactual way, that a certain causal process can change the value of a variable in order to examine the consequence on the other variable. In the present case, one can imagine to 'freeze' progressively the degrees of freedom, for example by constraining or selecting the particles.[4] As more and more particles are frozen, the ratio

---

[4]This is precisely the kind of intervention that only a Maxwell's Demon can carry out. However, as said above, it is not required that the intervention be physically possible. The possibility of acting in

of the equilibrium region to the other volumes shrinks. Finally, the relation governing this behavior depends only on the total energy of the system, therefore, "it would continue to hold—would remain stable or unchanged—as various other conditions change" (Woodward 2000, 205), which is Woodward's definition of invariant relationship. Thus, the number of degrees of freedom satisfies the conditions for being a causal factor in the manipulative sense: it allows us to control in an invariant way those qualitative characteristics that feature in our experience of equilibrium.[5]

Thus, (T1) conveys a sort of causal information that can be used for explanatory purpose. However, although the high number of degrees of freedom acts somewhat causally, it still does not look like 'the cause' of equilibrium. I will try to cope with this intuition with the following consideration. The huge difference in size between the equilibrium macrostate and the others yields very good reasons to believe that our actual experience will always be one in which systems converge toward equilibrium. The difference is so astronomically large that no reasonable agent could even expect anything different from equilibrium. The high number of degrees of freedom produces and influences the structure that realizes this expectation. However, the concrete realization in space and time of equilibrium also depends on the features of the system trajectory. The situation is similar to the case of a fair coin. A bit of combinatorics is sufficient to convince us that it is high unreasonable to expect 1,000,000 consecutive tails. However, we feel that the complete explanation of this fact has to do also with the way in which the coin is flipped. For this reason, in the typicality argument (T1) should be connected with the typicality-claim (T2). From this point of view, the explanatory information incapsulated in (T1) seems to be a species of Elliot Sober's equilibrium explanation

---

a "demonic" fashion is warranted by the laws of mechanics.

[5]Let me add a word of caution. The high number of degrees of freedom plays a causal role in the Boltzmannian approach subscribed by most typicalists. This approach considers only closed systems and therefore excludes external interventions. The situation is different with open systems.

(Sober 1983). This is, I suspect, the deep reason behind some of the above-mentioned criticisms leveled against the use of measure by typicalists. Statements about measure and about the structure of the phase space, such as (T1), still fall short of telling us a full causal story of the actual behavior of the system. To capture the difference between the causal explanatory information contained in (T1) and (T2) respectively, I introduce the difference between causal structure and causal scenarios. The number of degrees of freedom endows the phase space with a causal structure, i.e., the partition of macrostates. This structure has a causal status because it yields difference-making information that allows us to potentially control the occurrence of the explanandum. However, this causal structure only constrains the possibilities by making the explanandum highly expectable. For this structure to be wholly explanatory, it is necessary that the specific causal scenarios, i.e., the Hamiltonian of the system, be consistent with it. The notion of "consistency" between causal structure and causal scenarios is supposedly made more explicit by the property Γ, but, as we will see in a minute, this is still a mudded notion.

5.2   (T2) as typicality-1

As I stated in Section 3, the crucial question of the typicality approach is to determine the property Γ that allows one to conclude the thermodynamic-like behavior on the basis of the well-established phase space fact (T1). It is possible to find isolated statements that seem to suggest that there is no real question after all: the conclusion follows straightforwardly from (T1) regardless of any characterization of the Hamiltonian. Consider, for instance, this quote of Nino Zanghí:

> The convergence to equilibrium of natural phenomena is neither a consequence of a new physical law, nor the effect of an attractor in the microscopic dynamics. The equilibrium macrostate does not attract anything, systems, typically, "fall in it" because, in the phase space, the equilibrium

19

> macrostate occupies a region enormously bigger than the others. (Zanghì 2005, 170)

A similar thought can be found in this passage by Detlef Dürr:

> What is typicality? It is a notion for defining the smallness of sets of (mathematically inevitable) exceptions and thus permitting the formulation of law of large numbers type statements. Smallness is usually defined in terms of a measure. What determines the measure? In physics, the physical theory. Typicality is defined by a measure on the set of "initial conditions" (eventually by the initial conditions of the universe), determined, or at least strongly suggested by the physical law. Are typical events most likely to happen? No, they happen because they are typical. But are there also atypical events? Yes. They do not happen, because they are unlikely? No, because they are atypical. But in principle they could happen? Yes. So why don't they happen then? Because they are not typical. (Dürr 2001, 131)

These quotes do not represent the whole of the position of Zanghí and Dürr, but rather prove Frigg's point that typicalists have made only sporadic and sometimes contradictory remarks on how to interpret the dynamical part of the typicality argument (Frigg 2009). Be that as it may, the idea at times surfaced that one can conclude the argument from (T1) without introducing any specification of the Hamiltonian. I think that this strategy is epistemologically questionable for two reasons. First, it amounts to eluding the problem of typicality and reduces it to mere hand-waiving. Second, it leaves us with an incomplete explanation of equilibrium. As said, (T1) can at best picture a causal structure, but it does not produce more specific causal scenarios.

In the literature, one can find interpretations of (T2) as either typicality-1 or typicality-2. Let us begin with the former. A typicality-1 reading requires the introduction of a way of counting Hamiltonians that produce dynamical trajectories with a certain property. Recently, Roman Frigg and Charlotte Werndl have put forward an interesting attempt in this direction (Frigg and Werndl 2012, Werndl 2013). They begin with outlining the typicality argument presented in Section 3. With regards to this argument,

they set two questions. The first is "the conceptual question of whether we can explain the behavior of a particular system by appeal to what systems typically do" (Frigg and Werndl 2012, 920). This is essentially the epistemological issue of this paper. The other question concerns "what notions of typicality are at work in the two [typicality-claims]." By this, Frigg and Werndl mean something different from the semantic analysis carried out in this paper. They investigate what it means for an Hamiltonian to be typical-as-vast-majority. The obvious difficulty with this approach is that "function spaces, unlike phase spaces, do not come equipped with normalized measures" (Frigg and Werndl 2012, 921), hence their proposal amounts to replacing the measure-based notion of typicality-1 adopted for microstates, with a new topological notion. Frigg and Werndl introduce the concept of comeagre set as the topological equivalent of measure-1 set and define typicality as a relational property that an element of a set with a certain property P possesses if the set is comeagre respect to a suitable topology. The bulk of their paper is dedicated to show that the property Γ that Hamiltonian must have to produce thermodynamic-like behavior is epsilon-ergodicity, that the topology to use is the Whitney topology and that perturbed Lennard-Jones Hamiltonians fall into this category. The connection between (T1) and (T2) is ensured by epsilon-ergodicity. Basically, the trajectory will roam the phase space, therefore it will spend much more time in the equilibrium region and will tend to come back to it. The property of epsilon-ergodicity shows how the difference in size of the phase regions matters for the trajectory.[6]

This brilliant example of how to read (T2) in terms of typicality-1 suggests some important points on the explanatory value of typicality. While (T1) describes a causal structure compatible with multiple causal scenarios, (T2) aims at determining these

---

[6]To put it differently, Frigg and Werndl show that the vast majority of the Hamiltonians produce a trajectory in which the ratio between microstates of the different regions is very close to the ratio of their phase volumes. These trajectories have thermodynamic-like behavior.

causal scenarios.[7] By interpreting (T2) in terms of typicality-1, Frigg and Werndl proposed a deeper characterization of the family of causal scenarios underlying equilibrium. By establishing epsilon-ergodic Hamiltonian, they determine a what-if-things-had-been-different relation between microscopic variables and the causal structure of the phase space. In particular, their proposal communicates information about the general features of the causal scenarios.

5.3 (T2) as typicality-2

Although typicality-1 remains the most popular way to look at typicality, physicists have also alluded, albeit sometimes unwittingly, to an interpretation of (T2) in terms of typicality-2. The main thought of this approach is that the sought-after property Γ is a feature of the best exemplar of the trajectory produced by the underlying dynamics. The previous discussion of typicality-2 in Lewis' theory of counterfactuals will help us clarify this line of argument. In the literature, we can find two main ways to pursue this option. The first way consists in arguing that, given (T1), a dynamics that would not lead to thermodynamic-like behavior would be a sort of monstrosity. This is the claim expressed, for instance, by Sheldon Goldstein:

> [The phase space] consists almost entirely of phase points in the equilibrium microstate [...] with ridiculously few exceptions whose totality has volume of the order of $10^{-10^{20}}$ relative to that of the [phase space]. For a nonequilibrium phase point $[x]$ of energy E, the Hamiltonian dynamics governing the motion $[x_t]$ arising from $[x]$ would have to be ridiculously special to avoid reasonably quickly carrying $[x_t]$ into [the equilibrium state and keeping it there for an extremely long time—unless, of course, $[x]$ itself were

---

[7]The concept of causal scenario is close to the concept of "causal pattern" recently introduced by Angela Potochnik to extend equilibrium explanations (Potochnik 2015). A causal pattern is less specific than a causal process as introduced by Wesley Salmon and Phil Dowe, but it is more precise than Sober's causal structure. The explanatory value of causal patterns hinges on two factors: "(1) they feature one of more of the property of a system upon which the phenomenon to be explained depends and (2) they communicate information about the scope of that dependence" (Potochnik 2015, 1169).

ridiculously special. (Goldstein 2001, 43-44)

The grounding intuition of this argument is that a trajectory that stubbornly remains confined to a sequence of non-equilibrium states would be an extremely remarkable event, a fantastically unlikely concours of circumstances. A more handy way to express the same idea is phrased by the metaphor of the "conspiracy":

> For the macroscopic systems we are considering, the disparity between relative sizes of the comparable regions in the phase space is unimaginably larger. The behavior of such systems will therefore be as observed in the absence of any 'grand conspiracy'. (Lebowitz 1993b, 9)

The point of this argument is that, given the dominance of the equilibrium state, the only way to avoid the system to reach equilibrium is to intervene specifically on the microscopic trajectory in order to confine it purportedly within the non-equilibrium regions. The idea that an external intervention is necessary to make the system behave 'anti-thermodynamically' pervades the Boltzmannian tradition and, more interestingly, also Lewis' approach. Boltzmann uses this idea to argue that a regular behavior cannot emerge spontaneously, but requires an active intervention, such as in the case of a perfect reversal of positions and velocities of the molecules:

> If we choose the initial configuration on the basis of a previous calculation of the path of each molecule so as to violate intentionally the laws of probability, then of course we can construct a persistent regularity. (Boltzmann 1896, I, 22)

A strikingly similar language appears in Lewis' treatment of counterfactuals. As we have seen, if nature is indeterministic, we can construe counterfactuals with an actual consequent and a made-up antecedent, which however are true in a possible world close to ours. To block this possibility, Lewis introduces the notion of "quasi-miracle", i.e. events which do not blatantly violate the laws of physics, but still do not belong to a world close to ours. Lewis' comments on this concept are particularly interesting:

The quasi-miracle would be such a remarkable coincidence that it would be quite unlike the goings-on we take to be typical of our world. Like a big genuine miracle, it makes a tremendous difference from our world. [...] My point is not that quasi-miracles detract from similarity because they are so very improbable. They are; but ever so many unremarkable things that actually happen, and ever so many other things that might happen under various counterfactual suppositions, are likewise very improbable. What makes a quasi-miracle is not improbability *per se*, but rather the remarkable way in which the chance outcomes seem to conspire to produce a pattern. (Lewis 1979, 60)

The similarity between Lewis's language and the thoughts expressed by Boltzmann, Goldstein, and Lebowitz is striking. In particular, they all insist on the notion of "non-conspiracy" and Goldstein's property of "ridiculously special Hamiltonian" seems to mean the same as Lewis' "remarkable coincidence".[8] According to this way of arguing, Γ is the property of being 'nonconspiratorial' or 'unremarkable'. Thus, one could read (T2) as saying that the Hamiltonian produces, as best exemplar, nonconspiratorial or unremarkable trajectories. By this, one means that these trajectories do not remain confined into the tiny non-equilibrium region for an unreasonably long time—where the unreasonability depends on the ratio of these regions to the much bigger equilibrium region. From (T1) and this interpretation of (T2), the conclusion of the argument follows. Is this reading of (T2) epistemologically feasible? Does it supplement the explanatory contribution provided by the fact that the vast majority of microstates belong to equilibrium, i.e., (T1)? This is Pitowsky's challenge: what does typicality offer us beyond measure? I argue that neither unremarkableness, nor non-conspiracy provide us with illuminating explanatory information.

To show that, let us begin with recalling that neither concept is grounded on low probability, but rather in some other circumstances: typical events happen often be-

---

[8]Note that the starting point of Frigg and Werndl is precisely the interpretation of "ridiculously special Hamiltonian" in terms of typicality-1.

cause they are typical, not the other way round, as Dürr put it. In the quote above, Lewis makes precisely the same point: it's the fact that events are unremarkable and non-conspiratorial that makes them occur frequently. Now, the concept of "non-conspiracy" seemingly refers to some objective feature of the Hamiltonian. However, the concept is uninformative because is circular. The non-conspiratorial dynamics is defined in no better way than a dynamics that allows us to use the dominance of the equilibrium (T1) to explain our experience. This concept makes (T2) parasitic on (T1) and does not add further explanatory information. Another way to see the same point is the following. The concept of non-conspiracy does not specify the dynamics in a way that can be used for control or for answering a what-if-things-had-be-different question. It commits us to the following description of dynamics:

(5)   The Hamiltonian is such that it is not constrained to remain confined
       into a specific region of the phase space.

But, as Woodward has argued, statements like (5), even when they are laws of nature, does not causally illuminate their objects. His example is the statement 'No material object can be accelerated from a velocity less than that of light to a velocity greater than that of light' (Woodward 2000, 206). Such a statement tells us something about the way the universe is constituted, but it does not give us any direct way to manipulate or control relations between variables. The status of (5) is totally analogous: it informs us that weird patterns do not appear in nature, but does not instruct us how to control them. Hence, the concept of non-conspiracy is objective, but circular, therefore uninformative.

By contrast, the concept of "remarkableness" is non-circular, but is informative only of our epistemic status vis-a-vis the system. As Williams has noted, claims about remarkableness seem to refer to the fact that quasi-miracles are remarkable events for

any well-informed rational agent. A trajectory stubbornly confined within the tiny non-equilibrium region would be analogous to 1,000 consecutive heads of a fair coin, a series of uninterrupted green lights when driving along Park Avenue, or water and wind forming footsteps on the sand: simply too good to be true. However, this does not give us information about the structure of the universe, but rather on our epistemic status. Put in other terms, remarkableness is related to *inference*, for example the design inference. It is the epistemic symptom of some causal scenario—and we can ground on it some inferences—but does not provide us with any specific information about this scenario.

However, in the literature one can find another way to read (T2) as a typicality-2 claim. This reading relies on the notion of chaoticity of the microscopic trajectories. As it is often the case, some hints to this interpretation can be already found in Boltzmann's work. As early as 1868, Boltzmann characterized 'atypical' behavior as unstable (Boltzmann 1868). He noticed that it is possible to construe special situations in which a gas will remain indefinitely in a state of non-equilibrium (e.g., by perfectly aligning the particles on a straight line), but these situations are easy to disrupt. The basic idea is that the underlying dynamics produces as best exemplars chaotic trajectories. The sense of remarkableness and contrivance attached to atypical trajectories is due to their intrinsic instability.

This idea featured in the debate on the irreversibility in 1894-95,[9] and it is still prominent among the 'typicalists'. Any external intervention that makes the dynamics "conspiratorial", such as the reversal of molecular velocities, results in very unstable trajectories. In practice, conspiratorial intervention requires a "perfect aiming" to drive the trajectory of the system into special regions and this trajectory is easily destroyed by the smallest perturbation. Lebowitz expresses this point by means of a vivid metaphor:

---

[9]For the details see (Badino 2015, 65-71).

> The situation is analogous to pinball-machine-type puzzles where one is supposed to get a small metal ball into a particular small region. You have to do things just right to get it in, but almost anything you do gets it out into larger regions. For the macroscopic systems we are considering, the disparity between the sizes of the comparable regions of the phase space is unimaginably larger. (Lebowitz 1993a, 37)

The perfect aiming of atypical trajectory is analogous to riding a bicycle backward: it is mechanically possible, but it is very difficult in practice, because the motion is unstable. Hence, the argument is that the Hamiltonian produces chaotic trajectories very sensitive to the smallest perturbation. To constrain a trajectory within a special behavior requires a very unstable perfect aiming because any perturbation would move the trajectory away from that region and back into the enormous equilibrium region. The property of chaoticity characterizes the idea of non-conspiracy in a much more substantial way. It provides information on the relevant Hamiltonians and therefore contributes to determine the family of causal scenarios responsible for equilibrium. Thus, the idea of chaoticity is certainly more epistemologically promising, although it still needs to be developed from a physical point of view.

6   Conclusion

Let me briefly summarize the main points of my discussion of the debate on typicality. The argumentative pattern relies on two typicality-claims (T1) and (T2). To make the argument conclusive, one has to characterize the property Γ so as to use the solidly grounded (T1) to obtain the conclusion. There are two ways to spell out a typicality-claim: typicality-as-vast-majority, or typicality-1, and typicality-as-best-exemplar, or typicality-2. In statistical mechanics, (T1) has been undisputedly interpreted as typicality-1. From an explanatory point of view, it provides the phase space with a causal structure in which the high number of degrees of freedom plays the role of causal factor.

The second typicality-claim has been interpreted either way. Frigg and Werndl have shown that it is possible to count the Hamiltonians for which the property Γ is epsilon-ergodicity. I have argued that this strategy informs us about the features of possible causal patterns responsible of equilibrium. One can also find passages in which Γ is taken to be non-conspiratorial/unremarkable or chaotic. While the former way of interpreting (T2) is circular or epistemic, the latter has the potentialities to convey information on the causal scenarios.

How does this discussion respond to our initial questions? As the reader may recall, the first question concerned the explanatory surplus of typicality with regard to measure. From my discussion, it follows that the impression that typicality does not add anything to measure stems from the common attitude to adopt a typicality-1 reading. As this is the usual way in which (T1) is interpreted and as (T1) plays a crucial role in the argument, the suspicion can raise that typicality means just measure-close-to-1. However, a typicality-claim can also be interpreted as typicality-2 and this meaning is not directly reducible to measure. This answer brings us to the second question concerning the kind of explanation offered by typicality. In the case of statistical mechanics, the argument requires the cooperation of two typicality-claims, which, I have argued, communicate two different types of causal information. While (T1) refers to the causal structure induced on the phase space by the high number of degrees of freedom, (T2) aims at specifying the causal scenarios determining the actual spacetime behavior. However, the determination of the proper causal scenarios requires the precise characterization of the property Γ, a goal that looks still out of our reach.

References

Badino, M.: 2015, *The Bumpy Road: Max Planck from Radiation Theory to the Quantum, 1896-1906*, Springer, New York.

Boltzmann, L.: 1868, Studien über das Gleichgewicht der lebendigen Kraft zwischen bewegten materiellen Punkten, *Wiener Berichte* **58**, 517–560.

Boltzmann, L.: 1896, *Vorlesungen über Gastheorie*, Barth, Leipzig.

Dodd, D.: 2011, Quasi-miracles, Typicality, and Counterfactuals, *Synthese* **179**, 351–360.

Dürr, D.: 2001, Bohmian Mechanics, *in* J. Bricmont, D. Dürr, M. C. Galavotti, G. C. Ghirardi, F. Petruccione and N. Zanghì (eds), *Chance in Physics: Foundations and Perspectives*, Springer, Berlin, pp. 115–132.

Dürr, D., Goldstein, S. and Zanghì, N.: 1992, Quantum Equilibrium and the Origin of Absolute Uncertainty, *Journal of Statistical Physics* **67**, 843–907.

Elga, A.: 2004, Infinitesimal Chances and the Laws of Nature, *Australasian Journal of Philosophy* **82**, 67–76.

Frigg, R.: 2007, Why Typicality Does Not Explain the Approach to Equilibrium, *in* M. Suarez (ed.), *Probabilities, Causes, and Propensities in Physics*, Springer, Berlin, pp. 77–93.

Frigg, R.: 2009, Typicality and the Approach to Equilibrium in Boltzmannian Statistical Mechanics, *Philosophy of Science* **76**(5), 997–1008.

Frigg, R. and Werndl, C.: 2012, Demystifying Typicality, *Philosophy of Science* **79**, 917–929.

Gaifman, H. and Snir, M.: 1980, Probabilities Over Rich Languages, Testing and Randomness, *Journal of Symbolic Logic* **47**(3), 495–548.

Goldstein, S.: 2001, Boltzmann's Approach to Statistical Mechanics, *in* J. Bricmont, D. Dürr, M. C. Galavotti, G. C. Ghirardi, F. Petruccione and N. Zanghì (eds), *Chance in Physics: Foundations and Perspectives*, Springer, Berlin, pp. 39–54.

Goldstein, S.: 2012, Typicality and the Notions of Probability in Physics, *in* Y. Ben-Menahem and M. Hemmo (eds), *Probability in Physics*, Springer, Dordrecht.

Hemmo, M. and Shenker, O.: 2012, *The Road to Maxwell's Demon. Conceptual Foundations of Statistical Mechanics*, Cambridge University Press, Cambridge.

Lebowitz, J. L.: 1993a, Boltzmann's Entropy and Time's Arrow, *Physics Today* **9**, 32–38.

Lebowitz, J. L.: 1993b, Macroscopic Laws, Microscopic Dynamics, Time's Arrow and Boltzmann's Entropy, *Physica A* **194**, 1–27.

Lebowitz, J. L.: 1999, Statistical Mechanics: A Selective Review of Two Central Issues, *Review of Modern Physics* **71**, 346–357.

Lewis, D.: 1973, *Counterfactuals*, Blackwell, Oxford.

Lewis, D.: 1979, Counterfactual Dependence and Time's Arrow, *in* D. Lewis (ed.), *Philosophical Papers*, Vol. II, Oxford University Press, pp. 32–51.

Lewis, D.: 1994, Humean Supervenience Debugged, *Mind* **103**, 473–490.

Pitowsky, I.: 2012, Typicality and the Role of the Lebesgue Measure in Statistical Mechanics, *in* Y. Ben-Menahem and M. Hemmo (eds), *Probability in Physics*, Springer, Berlin, pp. 41–58.

Potochnik, A.: 2015, Causal Patterns and Adequate Explanations, *Philosophical Studies* **172**, 1163–1182.

Sober, E.: 1983, Equilibrium Explanation, *Philosophical Studies* **43**, 201–210.

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

Werndl, C.: 2013, Justifying Typicality Measures of Boltzmannian Statistical mechanics and Dynamical Systems, *Studies in History and Philosophy of Modern Physics* **44**, 470–479.

Williams, J. R. G.: 2008, Chances, Counterfactuals, and Similarity, *Philosophy and Phenomenological Research* **77**(2), 385–420.

Woodward, J.: 2000, Explanation and Invariance in the Special Sciences, *British Journal for the Philosophy of Science* **51**, 197–254.

Woodward, J.: 2003, *Making Things Happen. A Theory of Causal Explanation*, Oxford University Press, Oxford.

Zanghì, N.: 2005, I fondamenti concettuali dell'approccio statistico in fisica, *in* V. Allori, M. Dorato, F. Laudisa and N. Zanghì (eds), *La natura delle cose: Introduzione ai fondamenti e alla filosofia della fisica*, Carocci, Rome, pp. 139–227.