

TYPICAL: A THEORY OF TYPICALITY AND TYPICALITY EXPLANATION

By

ISAAC WILHELM

A thesis submitted to the

School of Graduate Studies

Rutgers, The State University of New Jersey

In partial fulfillment of the requirements

For the degree of

Master of Science

Graduate Program in Mathematics

Written under the direction of

Sheldon Goldstein

And approved by

New Brunswick, New Jersey

May, 2018

ABSTRACT OF THE THESIS
TYPICAL: A THEORY OF TYPICALITY AND TYPICALITY EXPLANATION
by ISAAC WILHELM

Thesis Director:
Sheldon Goldstein

Typicality is routinely invoked in science and mathematics: typical gases in a low-entropy macrostate evolve to a higher-entropy macrostate; typical quantum systems exhibit probabilistic behavior; typical realizations of percolation models with $p > \frac{1}{2}$ contain an infinite open cluster. And typicality facts like these back many explanations. But what is it for something to be typical? And how do typicality facts explain? In this paper, I analyze the notions of typical properties, typical objects, and typicality explanations. I show how typicality is used to explain mathematical, statistical mechanical, and quantum mechanical phenomena. Finally, I argue that typicality is distinct from probability.

Thanks to Michael Kiessling, Joel Lebowitz, and especially Shelly Goldstein for much helpful feedback and discussion.

Table of Contents

Abstract	ii
1 Introduction	1
2 Analyses of Typicality	3
3 Typicality Explanations	7
4 Typicality and Statistical Mechanics	13
4.1 Technical Preliminaries	13
4.2 Explaining Equilibrium	14
4.3 Explaining the Approach to Equilibrium	15
5 Typicality and Bohmian Mechanics	18
6 The Difference Between Probability and Typicality	26
6.1 Formal Differences	26
6.2 Explanatory Differences	29
7 Conclusion	32

1 Introduction

Consider a large box that has been divided in two by a retractable barrier. A gas occupies the box's left half. Now suppose that the barrier is removed. At first, the gas remains on the left. But soon it expands, and after a little while, it is in the equilibrium macrostate: the state in which, from a macroscopic point of view, the gas does not appear to be changing. In other words, the gas's microstate—that is, the exact physical state of the gas, given by the positions and momenta of all its particles—evolves from the gas's initial macrostate to equilibrium.

Why is that? Because expanding gases are *typical*: nearly all microstates compatible with the gas's initial macrostate evolve to macroscopic equilibrium after a little while. In other words, typical gases evolve to equilibrium. Some do not, of course. For some initial microstates, the gas particles stay on one side of the container. But for nearly all initial microstates, gases with those microstates will eventually expand.

In this paper, I analyze typical properties and typical objects. I also provide a formalism for typicality explanations, which are explanations that cite facts about what is typical. Some formal accounts of typicality, applicable in some scientific contexts, have been proposed.¹ But there are no general accounts of what typicality is. And there is no account of how typicality explains. The present paper provides both.

There has been considerable resistance to the notion of typicality. And there has been considerable resistance to the claim that typicality facts can be explanatory. Frigg argues that contrary to Boltzmann (1897/2003b), Lebowitz (1993), Goldstein (2001), and others, typicality cannot be used to explain why gases evolve towards equilibrium (2011, p. 82). For the proffered explanations do not pay enough attention to the dynamics governing the gas. Uffink argues that typicality cannot be used to

¹For examples, see (Goldstein et al., 2010) and (Frigg & Werndl, 2012).

explain such thermodynamic phenomena because typicality explanations fall short of rigorous mathematical proof (2007, p. 980). Typicality statements cannot be explanatory, because the explanandum does not follow logically from the explanans. Sklar (1973) argues against one kind of typicality explanation, which says that to explain something, one sometimes need only show that the set of possible violations has very small measure. Sklar argues that even if the set of violations has small measure, it does not follow that those violations are unusual or exceptional (1973, p. 211).

Arguments like these gain traction because there is no theory of typicality and typicality explanations. If there were such a theory, resistance to typicality would probably be much less severe. Indeed, many of those who reject various typicality explanations are sympathetic to the underlying, intuitive idea. For example, along with Werndl, Frigg eventually came to argue in favor of one particular way of quantifying typicality (2012). So objections to typicality can be understood as requests for clarification. What is it for something to be typical? And how do typicality facts explain? Advocates of typicality have also raised questions like these.²

In this paper, I propose some answers to these questions. In Section 2, I analyze the notions of a typical property and a typical object. I also give several mathematical definitions of a notion which underlies these analyses: the notion of ‘nearly all’ objects of a particular sort. In Section 3, I propose a general formalism for typicality explanations, and I give several examples of mathematical explanations which invoke typicality facts. In Section 4, I present two typicality explanations of thermodynamic phenomena. In Section 5, I present the details of a typicality explanation in quantum mechanics. Finally, in Section 6, I discuss some differences between typicality and probability.

²For example, Goldstein writes that a comprehensive philosophical analysis of typicality explanations “would be most welcome” (2012, p. 70).

2 Analyses of Typicality

Pre-theoretically, something is typical just in case nearly all things, of a certain sort, are like it. Expanding gases are typical because nearly all possible initial microstates lead to expansion. Typical quantum systems exhibit probabilistic behavior because nearly all possible initial configurations do. So there is a close connection between something being typical and nearly all of some things being a certain way.

In this section, I make that close connection precise. I propose an analysis of typical properties and an analysis of typical objects. These are accounts of what typical properties and objects *are*. Collectively, they form an account of what typicality itself is. They precisify the pre-theoretic idea that something is typical just in case nearly all of the relevant things are a certain way.

Here is an analysis of typical properties.

TYPICAL PROPERTY

Let Γ be a set and let P be a property. P is *typical in* Γ if and only if ‘nearly all’ of the elements in Γ exemplify P .

For example, let Γ be the set of all possible microstates of a gas with some fixed energy, and let P be the property *is in equilibrium*. Then P is typical in Γ because nearly all of those microstates are in equilibrium (the details of this typicality fact are discussed in Section 4).

Different definitions of ‘nearly all’ yield different versions of TYPICAL PROPERTY. I will state some of those definitions soon. For now, here is an analysis of typical objects.

TYPICAL OBJECT

Let Γ be a set, let P be a property, and let x be a member of Γ . Then x is *typical_o* (relative to P and Γ) if and only if x has P and P is typical in Γ .

For example, let Γ be the set of all possible microstates of a gas with some fixed energy,

let P be the property *is in equilibrium*, and let x be the microstate of a particular gas which is in equilibrium. Then x is typical_o (relative to P and Γ) because x has P and P is typical in Γ . In other words, x is typical relative to being a gas in equilibrium.

Both of the above analyses invoke the notion of ‘nearly all’. That notion is the heart and soul of typicality: facts about typicality are just facts about what nearly all entities of a certain sort are like. I shall now present three of the most common definitions.

When Γ is finite, ‘nearly all’ can be quantified by counting. Let Γ be a large finite set, let P be a property, and let A_P be the set of elements in Γ which exemplify P . Then *nearly all* elements in Γ have P if and only if $\frac{|\Gamma \setminus A_P|}{|\Gamma|} \ll 1$. Call this the ‘counting-theoretic’ definition of ‘nearly all’.

In general, the sets Γ invoked in science and mathematics, for the purposes of formulating typicality facts, are infinite. So other definitions of ‘nearly all’ are required. A particularly simple one invokes cardinalities. Let Γ be an infinite set, let P be a property, and let A_P be the set of elements in Γ which exemplify P . Then ‘nearly all’ elements in Γ have P if and only if $|\Gamma \setminus A_P| < |\Gamma|$.³ Call this the ‘cardinality-theoretic’ definition of ‘nearly all’.

For example, on the cardinality-theoretic definition, nearly all real numbers are normal. Let Γ be the set of reals. A real number x is normal in base b just in case each sequence of digits of length n , in the base b expansion of x , appears with an average frequency of $\frac{1}{b^n}$. A real number is normal just in case it is normal in every base. Let P be the property of being normal; so A_P is the set of normal numbers. Then the set of elements in Γ which are not in A_P – the set of non-normal numbers – is countable. The cardinality of Γ is uncountable. So nearly all reals are normal.

There are cases in which Γ is infinite, but ‘nearly all’ should not be understood in terms of cardinals. For example, let Γ be a disk. Let S be a very small sector of

³One consequence of this definition is that if nearly all elements in Γ have P , then $|A_P| = |\Gamma|$.

Γ ; so S is shaped like a very thin slice of pie. Let P be the property of not lying within S , so that A_P is the set of all points which lie outside the pie slice. Then Γ has the same cardinality as the set of points in Γ that are not in A_P ; that is, Γ and S have the same cardinality. So on the cardinality-theoretic definition, it is not the case that nearly all points in the disk lie outside the small sector. Thus, the cardinality-theoretic definition of ‘nearly all’ is no good here. It provides an acceptable sufficient condition, but not an acceptable necessary condition, for nearly all of the elements in a set to exemplify a property.

A third definition of ‘nearly all’, used throughout statistical mechanics and quantum mechanics, gets around some such problems. In full detail, it is quite technically sophisticated. But the basic idea is extremely intuitive. Whereas the previous two definitions quantified ‘nearly all’ by count or cardinality, this one quantifies ‘nearly all’ with measures. Let Γ be a set, let P be a property, let A_P be the set of elements in Γ which exemplify P , and let m be a measure such that $0 < m(\Gamma) < \infty$. Then ‘nearly all’ of the elements in Γ exemplify P (relative to m) if and only if $\frac{m(\Gamma \setminus A_P)}{m(\Gamma)} \ll 1$. Call this the ‘measure-theoretic’ definition of ‘nearly all’.

For example, consider the disk Γ and the small sector S once more. As before, P is the property of lying outside S , and A_P is the set of points in Γ which exemplify P . Let m be the standard measure of the areas of two-dimensional shapes; m is called the two-dimensional Lebesgue measure. Then nearly all of the elements in Γ exemplify P (relative to m). To see why, note that the area of the set of points in Γ but not in A_P —the area of the set of points in S —is much smaller than the area of Γ . So $\frac{m(\Gamma \setminus A_P)}{m(\Gamma)} \ll 1$.

The analyses TYPICAL PROPERTY and TYPICAL OBJECT do not include a free parameter for a measure. So to use the measure-theoretic definition of ‘nearly all’, relativize both sides of the biconditional in each analysis to a measure parameter. The resulting version of TYPICAL PROPERTY is: P is *typical in* Γ (relative to m)

if and only if nearly all of the elements in Γ exemplify P (relative to m). And the resulting version of TYPICAL OBJECT is: x is *typical_o* (relative to P , Γ , and m) if and only if x has P and P is typical in Γ (relative to m).

This completes my analysis of typicality. Roughly put, typical properties are exemplified by nearly all members of the relevant set, and typical objects are objects which exemplify typical properties. The three definitions of ‘nearly all’ give rise to many different versions of these analyses. But in each case, the basic idea is the same. Typicality is nearly all.

3 Typicality Explanations

Now that precise analyses of typicality are on the table, it is clear that there are typicality facts. It is clear that typicality is a real feature of our world.

What is less clear, however, is that typicality facts *explain*. Many different explanations are backed by facts about what is typical. In this section, I propose a formalism for typicality explanations, and I provide several examples of explanations which invoke typicality facts.

The basic schema for typicality explanation is as follows.

$$\frac{x \in \Gamma \quad P \text{ is typical in } \Gamma \text{ (relative to } m\text{)}}{\therefore x \text{ has } P} \quad (1)$$

where x is a particular entity, Γ is a set, P is a property, and m is a measure. This is a schema because different substitutions for the terms x , Γ , P , m and the ‘nearly all’ in the analysis of typicality yield different explanations. If that ‘nearly all’ is not defined measure-theoretically, then the parenthetical in the second line should be dropped. The symbol \therefore indicates that the deduction is a typicality deduction, which is different from a logical deduction \therefore because in typicality deductions, the conclusion is not a logical consequence of the premises. And as usual for theories of explanation, for an instance of (1) to be explanatory, the premises and the conclusion must all be true.

The typicality fact in the second premise of (1) does the bulk of the explanatory work. For it says that nearly all elements of Γ are a certain way. So it is no wonder that x is that way too, given that x is in Γ . Thus, the heart of a typicality explanation is the typicality fact that it invokes.

Many different explanations are typicality explanations. In sections 4 and 5, I

discuss two scientific examples of typicality explanations in detail. In the remainder of this section, I present a few examples of typicality explanations in mathematics.

The first concerns the surface areas of spheres. Suppose a mathematician is studying a series of points on the surface of high-dimensional sphere which satisfy some given properties. She finds that every single one of these points happens to be very close to the sphere's equator. Her friend, who is also studying points on the surface of a high-dimensional sphere, but who is interested in a very different set of properties, has encountered the same phenomenon: all of the friend's points are also extremely close to the equator. This is surprising. Why should points, which are more-or-less randomly selected from the sphere's surface, all have this one feature in common? Why should they all be extremely close to the equator?

The answer is a typicality fact: for large n , nearly all points on the surface of an n -dimensional sphere are very close to the equator. This follows from a more general result concerning the measures of subsets of spheres. Let σ_n be the uniform probability measure on sphere $\mathbb{S}^n = \{x \in \mathbb{R}^{n+1} \mid \|x\| = 1\}$.⁴ For each measurable $A \subseteq \mathbb{S}^n$ and each $t \in [0, \infty)$, define $A_t = \{x \in \mathbb{S}^n \mid \text{dist}(x, A) < t\}$.⁵ The more general result is as follows (Ledoux, 2001, p. vii):

Theorem 1. *For each measurable $A \subseteq \mathbb{S}^n$ such that $\sigma_n(A) \geq \frac{1}{2}$, and for each $t > 0$, $1 - \sigma_n(A_t) \leq e^{-\frac{t^2(n-1)}{2}}$.*

In other words, for small t and for very large n , the measure of the set of points in \mathbb{S}^n but not in A_t is extremely small.

It follows that for suitably large n , nearly all points on the sphere are within ϵ of the equator, where 'nearly all' is defined measure-theoretically. To see why, let $\delta \ll 1$, let $t = \epsilon$, let A be the northern hemisphere of \mathbb{S}^n , and let B be the southern

⁴For all $x = (x_1, \dots, x_n) \in \mathbb{R}^n$, $\|x\| = \left(\sum_{i=1}^n x_i^2\right)^{\frac{1}{2}}$.

⁵For any set $A \subseteq \mathbb{S}^n$ and any $x \in \mathbb{S}^n$, $\text{dist}(x, A) = \inf_{y \in A} \|x - y\|$.

hemisphere. For sufficiently large n , $e^{-\frac{t^2(n-1)}{2}} < \frac{\delta}{2}$. By Theorem 1, it follows that $\sigma_n(\mathbb{S}^n \setminus A_t) < \frac{\delta}{2}$, since $\sigma_n(\mathbb{S}^n \setminus A_t) = 1 - \sigma_n(A_t)$. Similarly, Theorem 1 implies that $\sigma_n(\mathbb{S}^n \setminus B_t) < \frac{\delta}{2}$. Therefore, $\sigma_n(\mathbb{S}^n \setminus (A_t \cap B_t)) < \sigma_n(\mathbb{S}^n \setminus A_t) + \sigma_n(\mathbb{S}^n \setminus B_t) < \delta$. The set $A_t \cap B_t$ is the set of all points on the sphere which are within ϵ of the equator. So by the measure-theoretic definition of ‘nearly all’, nearly all points on the sphere are near the equator.

This typicality fact relieves the surprise one might feel upon discovering that a given point on the surface of a high-dimensional sphere is near the sphere’s equator. The surface area of a high-dimensional sphere is concentrated near the equator, so given any point on the sphere, it should be unsurprising if that point is near the equator too.⁶

Infinite graphs in percolation models provide another example of typicality explanation in mathematics. A graph is a pair $\langle V, E \rangle$ where V is a set of vertices, E is a set of edges, each edge connects exactly two vertices, and any two vertices have at most one edge between them. Let $V = \mathbb{Z}^2$, and fix $p \in [0, 1]$. For each pair of adjacent vertices $i, j \in \mathbb{Z}^2$, p is the probability that there is an edge connecting i and j . Whether or not there is an edge between any pair of vertices is independent of whether or not there is an edge between any other pair of vertices. This distribution over possible graphs is called a percolation model; denote it by \mathcal{L}_p .

Some realizations of \mathcal{L}_p —that is, some graphs on \mathbb{Z}^2 —have an infinite open cluster. Infinite open clusters are defined as follows. A subgraph of a graph $\langle V, E \rangle$ is a graph $\langle V', E' \rangle$ such that $V' \subseteq V$ and E' is the set of all edges in E between vertices in V' . A connected subgraph of a graph $\langle V, E \rangle$ is a subgraph $\langle V', E' \rangle$ such that for any pair of vertices $i, j \in V'$ there is a sequence of edges in E' connecting i

⁶The same result holds for any other geodesic on an n -dimensional sphere, when n is large. For any given geodesic, nearly all points on the surface of such a sphere are within ϵ of that geodesic. It does not follow, of course, that nearly all points on the surface of such a sphere are within ϵ of *every* geodesic. In statements about typicality, the order of the quantifiers is extremely important: even though, for each geodesic g , nearly all points on the surface of such a sphere are within ϵ of g , it is false that nearly all points on the surface of such a sphere are within ϵ of each g .

to j . An infinite connected subgraph of a graph is a subgraph with infinitely many edges. Finally, an infinite open cluster of a graph is an infinite connected subgraph of that graph.

Realizations of \mathcal{L}_p are investigated for a variety of reasons: they help us understand more physically realistic models of ferromagnets, such as Ising models (Grimmett, 1989), for instance. When $p > \frac{1}{2}$, finite sections of those realizations are often found to contain extremely large connected subgraphs. As the size of the finite section increases, the connected subgraph therein tends to get larger and larger, suggesting that the realization of \mathcal{L}_p in question contains an infinite open cluster. But when $p \leq \frac{1}{2}$, finite sections of those realizations are often found to lack large connected subgraphs. This suggests that the realization in question lacks any infinite open clusters. So why is that? Why do realizations of \mathcal{L}_p appear to contain infinite open clusters for $p > \frac{1}{2}$, but not for $p \leq \frac{1}{2}$?

The explanation, once again, is a typicality fact. It can be shown that nearly all realizations of \mathcal{L}_p for which $p > \frac{1}{2}$ contain exactly one infinite open cluster, and nearly all realizations of \mathcal{L}_p for which $p \leq \frac{1}{2}$ contain no infinite open cluster (Kesten, 1980, p. 42). So it should not be surprising that realizations of \mathcal{L}_p seem to contain an infinite open cluster when $p > \frac{1}{2}$, and to lack an infinite open cluster when $p \leq \frac{1}{2}$. Given that nearly all realizations of \mathcal{L}_p adhere to that pattern, it is to be expected.

The theory of finite graphs provides striking examples of typicality examples too. One concerns the sizes of complete subgraphs—called ‘cliques’—of large graphs. A complete graph is a graph such that every pair of vertices is connected by an edge. A clique is a set V of vertices in a graph G such that the induced subgraph of V —the graph whose vertices are the members of V and whose edges are the edges in G between vertices in V —is complete.

Suppose one is studying lots of different large finite graphs on n vertices, and suppose one finds that the maximal cliques of these graphs all have exactly the same

size. That is highly surprising. Why should these disparate graphs all have this feature in common?

It turns out that there is a simple typicality explanation of this surprising fact. It can be shown that for sufficiently large n , the largest cliques of nearly all graphs of size n have *exactly* the same size. Typically, the size of the maximal clique will be approximately $\frac{2 \ln(n)}{\ln(1/p)}$ (Alon & Spencer, 2000, p. 159). This typicality fact relieves the surprise one might feel if, given a lot of different large finite graphs on n vertices, one finds that all of those graphs have maximal cliques of the exact same size.

Typicality analyses are especially useful when it is difficult, or downright impossible, to prove that a specific system has some property. The example of large cliques in finite graphs is a case in point. Because it is generally impossible to identify the maximal clique of a given graph in polynomial time, “the problem of finding the largest clique in a general graph is intractable. It is thus natural to study this problem for appropriately randomly generated input graphs” (Alon et al., 1998, p. 458). So in addition to explaining otherwise surprising facts, typicality analyses often provide tractable approaches to problems for which no other tractable approaches are forthcoming.

For another example of the utility of typicality analyzes like these, suppose we wish to analyze the properties of the Hamiltonians of quantum systems, such as the densities of the eigenvalues of those Hamiltonians per unit energy interval (Wigner, 1967, p. 4). In many cases, the specific Hamiltonian of a system is unknown, or it is difficult to analyze on its own. Nevertheless, one can often discover whether or not a given statistical property is the same for nearly all Hamiltonians. If so, then it is reasonable to infer that the specific Hamiltonian at issue has that property too.

Wigner, in fact, makes precisely this point, in regards to the development of random matrix theory. He writes that in many cases,

“[o]ne...deals with a specific system, with its proper (though in many cases un-

known) Hamiltonian, yet pretends that one deals with a multitude of systems, all with their own Hamiltonians, and averages over the properties of these systems. Evidently, such a procedure can be meaningful only if it turns out that the properties in which one is interested are the same for the vast majority of the admissible Hamiltonians. The first question, then, is what are the admissible Hamiltonians, and what is the proper measure in the ensemble of these Hamiltonians. The second question is, of course, whether, given the ensemble of admissible Hamiltonians with a proper measure, the properties in which we are interested are common for the vast majority of them". (1967, p. 3)

In other words, one way to analyze the properties of a specific Hamiltonian is to analyze the properties of 'admissible' Hamiltonians; all Hamiltonians of a certain well-defined kind. If a property is shared by the vast majority of admissible Hamiltonians, then it is reasonable to infer that the specific Hamiltonian has that property too.

4 Typicality and Statistical Mechanics

In this section, I discuss typicality explanations of statistical mechanical phenomena. I focus on the example mentioned in Section 1, in which a gas occupies the left half of a large box, and is kept on the left by a retractable barrier. When the barrier is removed, the gas begins to expand, and after a little while, it is in equilibrium.

The final equilibrium state of the gas, and the gas's approach to equilibrium, can both be explained using the formalism in Section 3. Each explanation relies on the theory of statistical mechanics. In this section, I define the relevant statistical mechanical notions, and I present the two explanations.

4.1 Technical Preliminaries

Let N be the number of particles in the gas; suppose that N is at least 10^{20} . Let $\Lambda = \mathbb{R}^{6N}$ be the N -particle phase space. Each phase point $X \in \Lambda$ can be written $X = (\vec{q}, \vec{p})$, where $\vec{q} = (q_1, \dots, q_N)$, $\vec{p} = (p_1, \dots, p_N)$, and for each i , $q_i \in \mathbb{R}^3$ is the position of particle i and $p_i \in \mathbb{R}^3$ is the momentum of particle i .

The evolution of a phase point is given by the following equations:

$$\frac{d\vec{p}}{dt} = -\frac{\partial \mathcal{H}}{\partial \vec{q}} \quad (2)$$

$$\frac{d\vec{q}}{dt} = \frac{\partial \mathcal{H}}{\partial \vec{p}} \quad (3)$$

Here, \mathcal{H} , called the Hamiltonian, is a function $\mathcal{H} : \Lambda \rightarrow \mathbb{R}$ which maps each phase point X to a number representing the total energy of X . Solutions to (2) and (3) yield a function $\phi_t : \Lambda \rightarrow \Lambda$ such that $\phi_t(X)$ is the microstate to which X evolves in

time t .

Consider again the gas in the box. Let $X(t) \in \Lambda$ denote the microstate of the gas at time t . Let E be the energy of microstate $X(0)$. By conservation of energy (\mathcal{H} is a constant of the motion), E is the energy of each subsequent microstate $X(t)$ as well. Let Λ_E be the set of all microstates in Λ whose energy is E .

Λ_E can be partitioned into sets called macrostates. To define a macrostate, divide the box into K cells, where K is large but $K \ll N$. Specify, to within a given accuracy, the number of particles in each cell and the total energy of each cell (Lebowitz, 1999, p. S347).⁷ Each such specification singles out a collection of microstates in Λ_E : those microstates which agree with that specification. The resulting collection is a macrostate. In general, let $M(X)$ denote the macrostate to which microstate X belongs.

One macrostate in Λ_E is particularly important for explaining the evolution of gases: the equilibrium macrostate M_{eq} . The important feature of M_{eq} is its size: M_{eq} is, by far, the largest macrostate in Λ_E . The size of a set $A \subseteq \Lambda_E$, denoted $|A|$, is given by the modified Lebesgue measure. This measure is characterized by the fact that it is stationary: the modified Lebesgue measure of a set of microstates A is equal to the modified Lebesgue measure of the time-evolved set $\phi_t(A)$ of those microstates. It can be shown that for a gas with just 10^{20} particles, $|M_{eq}|$ is roughly $10^{10^{20}}$ times as big as all the other macrostates combined (Goldstein, 2001, p. 43). M_{eq} is huge.

4.2 Explaining Equilibrium

In this subsection, I present a typicality explanation of the fact that the gas is in equilibrium after a reasonably short amount of time t_f . Suppose that as a matter of empirical fact, t_f is an hour. Let $\Gamma = \Lambda_E$. Let P be the property of belonging

⁷There is an issue here, concerning how to define the total energy of a cell. See (Lebowitz, 1999) for discussion.

to M_{eq} at time t_f . Since M_{eq} is so massive, nearly all elements of Γ exemplify P (relative to the modified Lebesgue measure). Therefore, P is typical in Γ (relative to that measure).

This typicality fact can be used to explain why the gas is in equilibrium at time t_f .

$$\begin{array}{l} X(t_f) \in \Gamma \\ P \text{ is typical in } \Gamma \text{ (relative to the modified Lebesgue measure)} \\ \hline \therefore X(t_f) \text{ has } P \end{array} \quad (4)$$

In other words, the gas is in equilibrium at time t_f because its microstate is in Λ_E , and nearly all microstates in Λ_E are in equilibrium at time t_f .

This typicality explanation should remove any surprise one might feel to find the gas in equilibrium at time t_f . We should not be surprised, that is, that the macrovariable values of the gas are the same ones we eventually observed when we ran a similar experiment last week. Why? Because being in equilibrium is typical. The vast majority of phase space is taken up by Γ_E . So it is not surprising that the gas is in equilibrium.

4.3 Explaining the Approach to Equilibrium

In this subsection, I outline a typicality explanation of the gas's approach to equilibrium. This is what the typicality explanation of the gas's approach to equilibrium *would be*, if the typicality fact it invokes were proved in full generality (rather than merely proved in some cases).

Let t_f be as before. Let $t = 0$ be the time at which the retractable barrier is removed. Let $X(0)$ be the initial microstate of the gas, and let $M(X(0))$ be its macrostate at that time. Let $\Gamma = M(X(0))$. Let P be the property of evolving to

M_{eq} in time t_f ; so microstate X exemplifies P just in case $\phi_{t_f}(X) \in M_{eq}$. Let A_P be the set of points in Γ which exemplify P .

Here is the typicality explanation.

$$\begin{array}{l} X(0) \in \Gamma \\ P \text{ is typical in } \Gamma \text{ (relative to the modified Lebesgue measure on } \Gamma) \\ \hline \therefore X(0) \text{ has } P \end{array} \quad (5)$$

In other words, the gas evolves to the equilibrium macrostate by time t_f because its microstate is in $M(X(0))$, and nearly all microstates which begin in $M(X(0))$ are in M_{eq} at time t_f .

The first premise of (5) is true by definition. The conclusion is true because, as a matter of empirical fact, the gas is in M_{eq} at time t_f . The second premise is an open conjecture: it has not yet been proven in full generality. It certainly seems to be true, however. It is supported by arguments due to Boltzmann (1877/2015), and there are proofs of restricted versions of it in certain toy models.⁸

The literature features many explanations along the same lines as (5). While summarizing his famous H -theorem, for instance, Boltzmann writes: “[i]t is just for certain singular initial states that [equilibrium] is never reached, for example when all the molecules are initially moving in a line perpendicular to two sides of the container. For the overwhelming majority of initial conditions, on the other hand, [equilibrium is reached in a relatively short amount of time]” (1896/2003a).⁹ As Boltzmann says, a typicality fact about initial conditions explains why gases tend to approach equilibrium. Of course, there are exceptions to this general rule. Boltzmann’s example concerns a collection of particles, all confined to a single line which is perpendicular to two parallel sides of the container: these particles will forever bounce off of each other

⁸See (Lanford, 1974).

⁹For other examples of explanations along the lines of (5), see Lebowitz (1993; 1999), Goldstein (2001; 2012), and Lazarovici and Reichert (2015).

and the container's walls, and so never evolve to equilibrium. But that is a highly atypical situation. The overwhelming majority of microstates are not like that. The overwhelming majority of microstates approach equilibrium.

5 Typicality and Bohmian Mechanics

In the Bohmian interpretation of quantum mechanics, the Born rule can be explained via typicality. The explanation concerns two distributions: the distribution given by the Born rule, which states that the distribution of the results of what are called position measurements is $|\psi|^2$; and the empirical distribution of the actual positions of particles. Moreover, as I briefly discuss at the end of this section, this explanation has implications for the kind of knowledge which we can have of quantum systems.

Nearly a century of experiments have shown that the empirical distribution is always within experimental error of the Born rule distribution. On the Bohmian interpretation of quantum mechanics, this close connection between the two distributions can be explained. Typically, the empirical distribution for n qualitatively identical systems converges to the Born rule distribution as n goes to infinity. So in Bohmian mechanics, the Born rule is not a primitive theoretical posit. It can be explained by typicality facts.

In this section, I present that typicality explanation. I focus on a special case of Bohmian mechanics for a universe that has just one spatial dimension. I also focus on subsystems consisting of a single particle at a particular time. The more general explanation of the relationship between the empirical distribution and the Born rule distribution—for three spatial dimensions, and for subsystems consisting of any number of particles—is in all essentials the same. See (Dürr et al., 1992) for the more general treatment.

In Bohmian mechanics, the state of the universe at any time consists of two components: the configuration of the particles and the wavefunction at that time. Let $\Gamma = \mathbb{R}^N$ be the set of all possible particle configurations, where N is the number of particles in the universe. Each configuration $Q \in \Gamma$ can be written $Q = (Q^1, \dots, Q^N)$,

where for each i , $Q^i \in \mathbb{R}$ represents the position of particle i . For each $Q \in \Gamma$ and for each time t , $Q_t = (Q_t^1, \dots, Q_t^N)$ represents the configuration of all the particles in the universe at time t . The initial configuration, for example, is represented by Q_0 .

In addition, for each time t , let $\Psi_t(q) = \Psi(q_1, \dots, q_N)$ (for $q_i \in \mathbb{R}$) denote the universal wavefunction at t . Throughout this section, I assume that the universal wavefunction is normalized. That is, I assume that for each time t , $\int_{\Gamma} |\Psi_t(q)|^2 dq = 1$, where the integral is taken with respect to the standard Lebesgue measure on Γ .

Two equations govern the evolution of the particle configuration $q \in \Gamma$ and the wavefunction Ψ . The first, called the Schrödinger equation, is

$$i\hbar \frac{\partial}{\partial t} \Psi_t(q) = - \sum_{i=1}^N \frac{\hbar^2}{2m_i} \Delta_i \Psi_t(q) + V \Psi_t(q) \quad (6)$$

and the second, called the guidance equation, is the following first-order system of differential equations:

$$\frac{dQ^i}{dt} = \frac{\hbar}{m_i} \text{Im} \left(\frac{\frac{\partial}{\partial Q^i} \Psi_t(Q)}{\Psi_t(Q)} \right) \quad (7)$$

Note that there is one instance of (7) for each i from 1 to N . The parameter m_i in (6) and (7) represents the “mass” of particle i . The ‘Im’ function outputs the imaginary part of its input.

Together, equations (6) and (7) describe the evolution of the universe’s physical state. Note that this evolution is deterministic. For a given initial configuration Q_0 and a given Ψ_0 , the configuration Q_t and the wavefunction Ψ_t are determined for all later times.

The universal wavefunction can be used to define wavefunctions—called conditional wavefunctions—for subsystems. In general, whenever scientists perform calculations to predict the outcomes of a given experiment on a system of particles, they do not use the wavefunction of the entire universe. They assume that the system

of particles has a wavefunction of its own, and they perform their calculations using that wavefunction instead. Bohmian mechanics has the formal and physical resources to make that assumption precise. It allows for a rigorous definition of the wavefunction of a system of particles. That definition invokes (i) the positions of particles comprising the environment of that system, and (ii) the wavefunction of the universe.

To distinguish between the configuration of the subsystem and the configuration of its environment, I will use X to denote the configuration of the subsystem, and I will use Y to denote the configuration of the environment. Thus, we may write $Q = (X, Y)$ for the configuration of the universe. For the case of a single-particle subsystem, $X \in \mathbb{R}$ is the configuration of that subsystem, and $Y = (Y^1, \dots, Y^{N-1}) \in \mathbb{R}^{N-1}$ is the configuration of its environment. We may write $Q_t = (X_t, Y_t)$ for the configuration to which the configuration (X, Y) evolves in time t .

The conditional wavefunction $\psi_t(x)$ for a single-particle subsystem at time t is defined as follows:

$$\psi_t(x) = \Psi_t(x, Y_t) = \Psi_t(x, Y_t^1, \dots, Y_t^{N-1}) \quad (8)$$

In other words, the conditional wavefunction for the subsystem at time t is obtained by taking the actual positions of all other particles at t , and plugging those positions into the universal wavefunction at that time.¹⁰

Let me use the above technical notions to explain, more intuitively, what we want to show. Suppose that at time t , we do an experiment in which we take n qualitatively identical particles—all with the same conditional wavefunction—and measure their positions. For each region of space, we can use those measurements to compute the frequency with which particles are found in that region. The empirical distribution is the mathematical object which encodes those frequency facts. Of

¹⁰The conditional wavefunctions given in (8) need not be normalized. So whenever the conditional wavefunction is used in a probability formula, it is to be regarded as having been normalized, by dividing (8) by $(\int_{\mathbb{R}} |\Psi_t(x, Y)|^2 dx)^{\frac{1}{2}}$.

course, before even doing the experiment, we can use the Born rule distribution to compute the probability of finding a particle in that region. It turns out that for each such region, the frequency and the probability are approximately equal. That is what we want to explain. There is a close agreement between the frequencies given by the empirical distribution and the probabilities given by the Born rule distribution. But why is that? Why is the Born rule distribution the one that gets the empirical frequencies right?

Let us now say, more precisely, what the empirical distribution is. Take n one-particle subsystems: the one-particle subsystems corresponding to particles 1 through n . The empirical distribution at time t for the ensemble of particles is

$$\rho_{emp}(x) = \frac{1}{n} \sum_{i=1}^n \delta(x - X_t^i) \quad (9)$$

The empirical distribution is the probability measure which assigns an equal probability of $\frac{1}{n}$ to the position X_t^i of each particle (for $1 \leq i \leq n$). When the delta function $\delta(x - X_t^i)$ is integrated over a set $A \subseteq \mathbb{R}$, the result is 1 if $X_t^i \in A$ and 0 if $X_t^i \notin A$. So when the distribution $\rho_{emp}(x)$ is integrated over a set A , the result is the proportion of particles whose positions lie in A . That is how empirical frequencies are extracted from the empirical distribution: the frequency with which the particles lie in a given set A equals $\int_A \rho_{emp}(x) dx \equiv \rho_{emp}(A)$, the integral of ρ_{emp} over A .

As for the Born rule distribution, suppose that the particles all have the same conditional wavefunction relative to the chosen coordinate systems for those subsystems; denote it by ψ . Then the Born rule distribution is

$$\rho_{born}(x) = |\psi(x)|^2 \quad (10)$$

Thus, the Born rule probability for the event $X_i \in A$ is $\int_A \rho_{born}(x) dx \equiv \rho_{born}(A)$, the integral of ρ_{born} over A .

One finds, in practice, that for any given $A \subseteq \mathbb{R}$,

$$|\rho_{emp}(A) - \rho_{born}(A)| < \epsilon \quad (11)$$

for some positive $\epsilon \ll 1$. There is a typicality explanation for this fact, but it requires care to formulate properly. One reason for this is that most histories are such that consideration of (11) would be inappropriate, because for most histories, the conditional wavefunctions of the n one-particle subsystems are not all equal to each other. Another reason is that it could be the case that the information contained in the environment's configuration Y gives us more information about our subsystems, and that this alters the Born rule probabilities.

So we must show that this does not happen, and this means that we need a version of (11) that involves conditioning on the configuration Y_t of the environment at time t . Towards that end, let Ψ_0 be the initial wavefunction of the universe, and let Y be such that each conditional wavefunction of subsystem i is the same at time t ; so $\psi_t(x^1) = \psi_t(x^2) = \dots = \psi_t(x^n) \equiv \psi(x)$, where $\psi(x)$ is the common conditional wavefunction of each of the n one-particle systems. Define a distribution on initial configurations $Q_0 \in \Gamma$ by $P(dQ_0) = |\Psi_0(Q_0)|^2 dQ_0$. Then for time t , define the conditional distribution $P_t^Y(X_t \in dx) = P(X_t \in dx | Y_t = Y)$, where X_t is the random configuration of the composite n one-particle subsystems at time t and Y_t is the random configuration of the environment of that composite system at time t . The Bohmian equations of motion imply that the distribution of the random variable Q_t is equivariant: that is, with $P(dQ_0) = |\Psi_0(Q_0)|^2 dQ_0$, it follows that $P(dQ_t) = |\Psi_t(Q_t)|^2 dQ_t$ for all times t (Dürr et al., 1992, pp. 854-855). Therefore,

$$P(X_t \in dx | Y_t = Y) = |\psi_t(x_1, \dots, x_n)|^2 dx \quad (12)$$

where $\psi_t(x_1, \dots, x_n)$ is the conditional wavefunction of the composite subsystem con-

sisting of the n one-particle subsystems. Suppose, as one would expect from the fact that all the one-particle subsystems have the same conditional wavefunction $\psi(x)$, that at time t , the wavefunction of the n -particle system consisting of all one-particle subsystems equals the product of the wavefunctions for each of those systems individually.¹¹ That is, suppose that $\psi_t(x^1, \dots, x^n) = \psi(x^1) \cdots \psi(x^n)$. From (12), it then follows that $P(X_t \in dx | Y_t = Y) = P(X_t^1 \in dx_1, \dots, X_t^n \in dx_n | Y_t = Y) = |\psi(x^1)|^2 \cdots |\psi(x^n)|^2 dx^1 \cdots dx^n$. This implies that the X_t^i are independent.

Let B be the set of initial configurations $Q_0 = (X_0, Y_0)$ which evolve in such a way that at time t , the configuration of the environment Y_t is Y . Let G be the property of being an initial configuration such that (11) holds. For each i ($1 \leq i \leq n$), let $I_{\{X_t^i \in A\}}$ be the indicator function on A for X_t^i . Since the X_t^i are independent random variables, the $I_{\{X_t^i \in A\}}$ are independent random variables too. Note that

$$\begin{aligned} \rho_{emp}(A) &= \int_A \rho_{emp}(x) dx \\ &= \frac{1}{n} \int_A \sum_{i=1}^n \delta(x - X_t^i) dx \\ &= \frac{1}{n} \sum_{i=1}^n I_{\{X_t^i \in A\}} \end{aligned}$$

And therefore, by the law of large numbers, as the number of subsystems n gets large, $P_t^Y(G)$ —that is, the measure of the set of initial configurations which satisfy (11)—goes to 1. In other words, for large n , and for typical configurations in B , the n one-particle subsystems are distributed in accord with the Born rule distribution.

Now for the typicality explanation. Let $Q_0 \in \Gamma$ be the universe's initial configuration. The typicality explanation is as follows.

¹¹This can be derived from the effective wavefunctions for these subsystems; see (Dürr et al., 1992, pp. 866-867).

$$\begin{array}{c}
Q_0 \in B \\
G \text{ is typical in } B \text{ (relative to } P_t^Y) \\
\hline
\therefore Q_0 \text{ has } G
\end{array}
\tag{13}$$

Though lots of technical notions are involved in the formulation of (13), it is an intuitively satisfying explanation. Why does the empirical distribution of the n subsystems' actual configurations approximate, to arbitrarily high degree, the Born rule distribution for those systems? Because for nearly all configurations which result in the actual configuration of the environment, the empirical distribution of the n subsystems is extremely close to the Born rule distribution of those systems. So we explain the phenomenon of the Born rule by showing that among initial configurations which lead to the actual configuration of the environment, nearly all of them give rise to the Born rule distribution.

Of course, to simplify matters, I did not define everything needed to give a full explanation of the Born rule. I only discussed one-particle subsystems. I assumed that there is just one spatial dimension. And I focused on systems at a single time t . But none of these assumptions were necessary; I made them simply in order to make the technical definitions more accessible. Fully rigorous explanations of the Born rule, which are typicality explanations, can be found in (Dürr et al., 1992) and (Dürr & Teufel, 2009).

This typicality explanation shows that there is an absolute limitation on the knowledge of quantum systems that we can obtain. Even if we knew everything about the system's environment, we would not have any more information about the configuration of the system than the Born rule provides. For suppose we had more precise information about the system's configuration. That information would be reflected in correlations between the system and its environment: correlations between those features of the system to which the information refers, and those features of

the environment in which that information is represented. But given the system's wavefunction, there are no such correlations, which is more-or-less what (12) directly says. Even if we knew every microscopic detail of the environment, we would not know any more about the system's configuration than is given by the Born rule (Dürr et al., 1992, p. 883).

6 The Difference Between Probability and Typicality

The relationship between probability and typicality is subtle. Their close kinship, in my view, is articulated by the law of large numbers: according to my preferred reading, the law of large numbers says that the sample mean typically (rather than ‘probably’) approximates the population mean. But I shall not focus on that here.

Instead, I shall focus on the *differences* between probability and typicality. For in conversation, people often question whether typicality is anything over and above probability. What is typically the case, they tend to say, is just what is probably the case. And they often claim that typicality explanations are nothing over and above probabilistic explanations.

So it is worth spelling out the differences between probability and typicality in some detail. There are formal differences: not all typicality facts can be expressed using probability measures. And there are explanatory differences: some probabilistic explanations are not typicality explanations, and some typicality explanations are not probabilistic explanations.

6.1 Formal Differences

There are several big formal differences between probability and typicality. Here I focus on one: in a precise sense, typicality facts ‘outstrip’ probabilistic facts. Typicality is strictly more expressive, in the sense that some typicality facts—in which ‘nearly all’ is defined using cardinality—cannot be expressed using only probability measures.

Against this, one might claim that there is a probability measure m such that for each set Γ and each property P , if P is typical in Γ on the cardinality-theoretic definition of ‘nearly all’ then according to m , the set of elements in Γ which do not

exemplify P has extremely small probability. If this were true, then all typicality facts which rely on the cardinality-theoretic definition of ‘nearly all’ could be expressed via probability facts. But it can be shown that this claim is false. Here is a mathematically precise version of the claim: there exists a probability measure m such that for each set Γ and each property P , if $|\Gamma \setminus A_P| < |\Gamma|$ then $m(\Gamma \setminus A_P) \ll 1$ (where A_P is the set of elements in Γ which exemplify P). Suppose such an m exists. Since m is a function, m is defined over a particular domain; call it D . Let Γ be $\mathcal{P}(D)$, the power set of D . Then m is not defined over Γ , so $m(\Gamma \setminus A_P)$ is undefined for any $A \subseteq \Gamma$.

Alternatively, one might claim that for each Γ there is a probability measure m such that for each P , if P is typical in Γ on the cardinality-theoretic definition of ‘nearly all’ then according to m , the set of elements in Γ which do not exemplify P has extremely small probability. Again, if this were true, then all typicality facts which rely on the cardinality-theoretic definition of ‘nearly all’ could be expressed via probability facts. But again, it can be shown that this claim is false. Here is a mathematically precise version of the claim: for each set Γ there is a probability measure m over Γ such that for each property P , if $|\Gamma \setminus A_P| < |\Gamma|$ then $m(\Gamma \setminus A_P) < \epsilon$ for some $\epsilon > 0$ such that $\epsilon \ll 1$ (where A_P is the set of elements in Γ which exemplify P). To see that this is false, let $\Gamma = \mathbb{N}$, the set of natural numbers. For each $i \in \mathbb{N}$, let P_i be the property of being greater than i ; so $A_{P_i} = \{i + 1, i + 2, \dots\}$. Suppose for reductio that the claim holds. Note that for each i , $|\Gamma \setminus A_{P_i}| < |\Gamma|$. So by the claim, there exists a probability measure m such that for all i , $m(\Gamma \setminus A_{P_i}) < \epsilon$ (for some $\epsilon > 0$ such that $\epsilon \ll 1$). But by the upward continuity of measures,

$$\begin{aligned}
 1 &= m(\mathbb{N}) \\
 &= m\left(\bigcup_{i=1}^{\infty} \Gamma \setminus A_{P_i}\right) \\
 &= \lim_{i \rightarrow \infty} m(\Gamma \setminus A_{P_i})
 \end{aligned}$$

$$\begin{aligned}
&< \lim_{i \rightarrow \infty} \epsilon \\
&= \epsilon
\end{aligned}$$

which is a contradiction.

These technical issues reveal a significant formal difference between probability and typicality. While probability measures are upwards continuous, typicality is not. It follows that when ‘nearly all’ is defined cardinality-theoretically, the measures of certain typical sets must get arbitrarily small. So even though those sets are typical, their probabilities must get arbitrarily low.

Moreover, this situation is not particular to the cardinality-theoretic definition of ‘nearly all’. The exact same situation arises for other definitions. For example, consider the following ‘topology-theoretic’ definition of ‘nearly all’: ‘nearly all’ of the elements in Γ exemplify P (relative to topology τ) if and only if $\Gamma \setminus A_P$ is nowhere dense, where A_P is the set of elements in Γ which exemplify P .¹² On parallel with the previous discussion, one might claim that for each Γ and each topology τ on Γ , there is a probability measure m such that for each property P , if $\Gamma \setminus A_P$ is nowhere dense then $m(\Gamma \setminus A_P) < \epsilon$ for some $\epsilon > 0$ such that $\epsilon \ll 1$ (where A_P is the set of elements in Γ which exemplify P). But this is false. Once again, let $\Gamma = \mathbb{N}$. Let τ be the cofinite topology on Γ : so the open subsets of Γ are the sets whose complement is finite. As before, let P_i be the property of being greater than i ; so $A_{P_i} = \{i + 1, i + 2, \dots\}$. It can be shown that each $\Gamma \setminus A_{P_i}$ is nowhere dense. Suppose for reductio that the claim holds. Then there is a probability measure m such that for all i , $m(\Gamma \setminus A_{P_i}) < \epsilon$ (for some $\epsilon > 0$ such that $\epsilon \ll 1$). But just as before, the upward continuity of measures implies that $1 = m(\mathbb{N}) = \lim_{i \rightarrow \infty} m(\Gamma \setminus A_{P_i}) < \epsilon$, which is a contradiction.

¹²This definition of ‘nearly all’ is drawn from the discussion in (Frigg & Werndl, 2012).

6.2 Explanatory Differences

Probabilistic explanations and typicality explanations come apart. In particular, some probabilistic explanations are clearly not typicality explanations, and some typicality explanations are clearly not probabilistic explanations. That is another reason to think that typicality and probability are not the same. The explanations they support are sometimes different.

Here is an example of a probabilistic explanation which is not a typicality explanation, adapted from an example due to Railton (1978, p. 214). Let i be a uranium atom, let F be the predicate ‘is a uranium atom’, and let G be the predicate ‘emits an alpha particle’. Let r be the probability that a uranium atom undergoes such an emission within a particular interval of time. For the sake of the example, suppose that r is fairly small: $\frac{1}{2}$ or less, say. And suppose that i actually does emit an alpha particle. Now plug all these facts into Hempel’s schema for probabilistic explanation (1965).

$$\begin{array}{c} Fi \\ p(G, F) = r \\ \hline \therefore [r] \quad Gi \end{array} \tag{14}$$

Nowadays, it is generally assumed that some events can be explained by probabilities which are not high, and even by probabilities which are relatively low. So (14) is a probabilistic explanation.¹³

(14) is not a typicality explanation, however. Since r is relatively low, it is not the case that typical uranium atoms emit alpha particles in the time interval in question. For in that time interval, most do not. The property of emitting an alpha particle in the relevant time interval is not typical in the set Γ of uranium atoms.

¹³For discussion of the view that low probabilities can explain, see Salmon (1971, p. 56), Scriven (1959, p. 480), and Strevens (2000, p. 368).

So there is no typicality explanation of Gi . The second premise of the corresponding instance of explanatory schema (1)—the premise which would state that P is typical in Γ —is false in this case.

There is a typicality explanation which is clearly not a probabilistic explanation. It crops up within a version of the Everettian interpretation of quantum mechanics. According to this version of the Everettian interpretation, there is a multiplicity of approximately classical, approximately non-interacting regions of the wavefunction which can be described as classical worlds (Wallace, 2012, p. 38). These regions are often called ‘branches’, and together they comprise the Everettian multiverse.

Everettian quantum mechanics differs from orthodox quantum mechanics in many ways, but one will prove especially important here. Of course, in both orthodox quantum mechanics and Everettian quantum mechanics, there are multiple possible outcomes for any given experiment. In orthodox quantum mechanics, only one of those outcomes ever actually occurs after measurement. Only one of the many possible outcomes ultimately obtains. In Everettian quantum mechanics, however, all possible outcomes obtain after measurement. All are actual. For example, suppose an electron’s wavefunction is in a superposition of the electron being on the left and the electron being on the right. Suppose we do an experiment to detect an electron’s location, and suppose we find it on the left. In orthodox quantum mechanics, there is no electron on the right. Physical reality does not include the right-located electron. In Everettian quantum mechanics, however, there is an electron on the right. Each possibility—the electron being on the left, and the electron being on the right—corresponds to a distinct, and actual, physical situation. Each possibility obtains. We happened to detect the electron on the left, but the other electron is still there. It is just on a different branch.

In this version of Everettian quantum mechanics, there is a typicality explanation which cannot be a probabilistic explanation. The explanandum is that the

observed outcomes of our experiments, taken together, match the probabilities for those outcomes which the Born rule predicts. To explain this agreement between observation and the Born rule, let Γ be the set of all sequences of observations which we could have made. Let o be the sequence of observations which we did in fact make. Let P be the property of matching, to within some very small δ , the Born rule probabilities. As a matter of fact, o exemplifies P and o is in Γ . It can be shown that relative to a particular typicality measure R , P is typical in Γ (Barrett, 2017, pp. 33-35). So here is a typicality explanation of the fact that the observed sequence of outcomes conforms to the probabilities given by the Born rule.

$$\begin{array}{c} o \in \Gamma \\ P \text{ is typical in } \Gamma \text{ (relative to } R) \\ \hline \therefore o \text{ has } P \end{array} \quad (15)$$

In other words, our observations agree with the Born rule probabilities because the property of agreeing with those probabilities is typical in the set of all possible sequences (relative to R).

Clearly, (15) is a typicality explanation. But it is not a probabilistic explanation. For as discussed earlier, in this version of Everettian quantum mechanics, the various possible outcomes of any given experiment all obtain. Everett himself makes this point: it would be a mistake, he says, to think of just one outcome as obtaining, to the exclusion of the rest (1956/2012, p. 149). But in probabilistic explanations, *the explanandum is the only outcome, of the various possible outcomes, that occurs*. (15) cannot be a probabilistic explanation because the alternative possibilities—such as the possible sequences of observations which contradict the Born rule probabilities—occur too. Probabilistic explanations presuppose that only one of the mutually exclusive possibilities obtains: that is just what ‘mutually exclusive’ means. So (15) cannot be a probabilistic explanation.

7 Conclusion

Typicality is ‘nearly all’: typical properties are properties exemplified by nearly all elements of the relevant set, and typical objects exemplify typical properties. Typicality facts can back explanations. And typicality explanations are quite common: they appear throughout science and mathematics.

References

- Alon, N., Krivelevich, M., & Sudakov, B. (1998). Finding a Large Hidden Clique in a Random Graph. *Random Structures and Algorithms*, 13, (3-4), 457-466.
- Alon, N., & Spencer, J. H. (2000). *The Probabilistic Method* (2nd ed.). New York, NY: Wiley.
- Barrett, J. A. (2017). Typical worlds. *Studies in History and Philosophy of Modern Physics*, 58, 31-40.
- Boltzmann, L. (2003a). Reply to Zermelo's Remarks on the Theory of Heat. In S. G. Brush (Ed.), *The Kinetic Theory of Gases* (pp. 392-402). London: Imperial College Press. (Original work published 1896).
- Boltzmann, L. (2003b). On Zermelo's Paper "On the Mechanical Explanation of Irreversible Processes". In S. G. Brush (Ed.), *The Kinetic Theory of Gases* (pp. 412-419). London: Imperial College Press. (Original work published 1897).
- Boltzmann, L. (2015). On the Relationship between the Second Fundamental Theorem of the Mechanical Theory of Heat and Probability Calculations Regarding the Conditions for Thermal Equilibrium. (K. Sharp & F. Matschinsky, Trans.). *Entropy*, 17(4), 1971-2009. (Original work published 1877).
- Dürr, D., Goldstein, S., & Zanghì, N. (1992). Quantum Equilibrium and the Origin of Absolute Uncertainty. *Journal of Statistical Physics*, 67(5/6), 843-907.
- Dürr, D. & Teufel, S. (2009). *Bohmian Mechanics*. Berlin: Springer-Verlag.
- Everett, H. (2012). The Theory of the Universal Wave Function. In J. Barrett & P. Byrne (Eds.), *The Everett Interpretation of Quantum Mechanics* (pp. 72-172). Princeton, NJ: Princeton University Press. (Original work published 1956).
- Frigg, R. (2011). Why Typicality Does Not Explain the Approach to Equilibrium. In M. Suárez (Ed.), *Probabilities, Causes and Propensities in Physics* (pp. 77-93). New York, NY: Springer.
- Frigg, R., & Werndl, C. (2012). Demystifying Typicality. *Philosophy of Science*, 79, 917-929.
- Grimmett, G. (1989). *Percolation*. New York, NY: Springer.
- Goldstein, S. (2001). Boltzmann's Approach to Statistical Mechanics. In J. Bricmont, D. Dürr, M. C. Galavotti, G. Ghirardi, F. Petruccione, & N. Zanghì (Eds.), *Chance in Physics* (pp. 39-54). Heidelberg: Springer.
- Goldstein, S. (2012). Typicality and Notions of Probability in Physics. In Y. Ben-Menahem & M. Hemmo (Eds.), *Probability in Physics* (pp. 59-71). New York, NY: Springer-Verlag.
- Goldstein, S., Lebowitz, J., Mastrodonato, C., Tumulka, R., & Zanghì, N. (2010). Normal Typicality and von Neumann's Quantum Ergodic Theorem. *Proceedings: Mathematical, Physical and Engineering Sciences*, 466(2123), 3203-3224.
- Hempel, C. G. (1965). Statistical Explanation. *Aspects of Scientific Explanation and other Essays in the Philosophy of Science*. New York, NY: Macmillan.
- Kesten, H. (1980). The Critical Probability of Bond Percolation on the Square Lattice Equals $1/2$. *Communications in Mathematical Physics*, 74, 41-59.
- Lanford, O. E. (1974). Time Evolution of Large Classical Systems. In J. Moser (Eds.), *Dynamical Systems, Theory and Application* (pp. 1-111). New York,

- NY: Springer.
- Lazarovici, D., & Reichert, P. (2015). Typicality, Irreversibility and the Status of Macroscopic Laws. *Erkenntnis*, 80, 689-716.
- Lebowitz, J. (1993). Macroscopic laws, microscopic dynamics, time's arrow and Boltzmann's entropy. *Physics A*, 194, 1-27.
- Lebowitz, J. (1999). Statistical mechanics: A selective review of two central issues. *Reviews of Modern Physics*, 71(2), S346-S357.
- Ledoux, M. (2001). *The Concentration of Measure Phenomenon*. Providence, RI: American Mathematical Society.
- Railton, P. (1978). A Deductive-Nomological Model of Probabilistic Explanation. *Philosophy of Science*, 45, 206-226.
- Salmon, W. C. (1971). Statistical Explanation. In W. C. Salmon (Ed.), *Statistical Explanation and Statistical Relevance* (pp. 29-87). Pittsburgh, PA: University of Pittsburgh Press.
- Sklar, L. (1973). Statistical Explanation and Ergodic Theory. *Philosophy of Science*, 40(2), 194-212.
- Scriven, M. (1959). Explanation and Prediction in Evolutionary Theory. *Science*, 130(3374), 477-482.
- Strevens, M. (2000). Do Large Probabilities Explain Better? *Philosophy of Science*, 67, 366-390.
- Wallace, D. (2012). *The Everett Interpretation*. Oxford: Oxford University Press.
- Wigner, E. P. (1967). Random Matrices in Physics. *SIAM Review*, 9(1), 1-23.
- Uffink, J. (2007). Compendium on the Foundations of Classical Statistical Physics. In J. Butterfield & J. Earman (Eds.), *Philosophy of Physics* (pp. 923-1074). Amsterdam: Elsevier.