Typical

A Theory of Typicality and Typicality Explanation

Isaac Wilhelm

Abstract

Typicality is routinely invoked in everyday contexts: bobcats are typically short-tailed; people are typically less than seven feet tall. Typicality is invoked in scientific contexts as well: typical gases expand; typical quantum systems exhibit probabilistic behavior. And typicality facts like these support many explanations, both quotidian and scientific. But what is it for something to be typical? And how do typicality facts explain? In this paper, I propose a general theory of typicality. I analyze the notion of a typical property. I provide a formalism for typicality explanations, and I give an account of why typicality explanations are explanatory. Along the way, I show how typicality facts explain a variety of phenomena, from everyday phenomena to the statistical mechanical behavior of gases. Finally, I argue that typicality is not the same thing as probability.

This…is…typical.

— Basil Fawlty

1 Introduction

Typical bobcats have short tails. Some do not, of course. Because of genetic mutations, some bobcats' tails are long. Because of cruel trapping practices, some bobcats have no tails at all. But nearly all bobcats have short tails. Having a short tail is typical among bobcats.

In this paper, I propose a theory of typicality. I analyze the notion of a typical property, and I offer an account of typicality explanations. As I argue, facts about typical properties are explanatory: we use them, quite often, to explain phenomena.

My theory is motivated by discussions in the philosophy of science literature: typicality is invoked in statistical mechanics to explain thermodynamic behavior, and it is invoked in quantum mechanics to explain the appearance of quantum probabilities. Some formal accounts of typicality, applicable in some scientific contexts, have been proposed.[1] But there are no general accounts of the metaphysics of typicality, of what typicality *is*. And there are no systematic accounts of how typicality explains.[2] This paper provides both.

There are other reasons to pursue a philosophical account of typicality. We use typicality all the time, in the most quotidian contexts. You go for a hike in the woods with a friend. A Pennsylvania bobcat crosses your path. Your friend asks 'Why does that creature have such a short tail?' 'Because it is a bobcat', you reply, 'and having a short tail is typical among bobcats'. This explanation is perfectly legitimate. Your inquiring friend is not a biologist;

[1]For example, see (Frigg and Werndl [2012]) and (Goldstein [2001]).

[2]For some discussion of typicality and explanation, see (Lazarovici and Reichert [2015]) and (Maudlin [2011]).

she is not asking about the genetic make-up of the bobcat. Nor is your inquiring friend a mereologist; she is not interested in the part-whole relation that obtains between the tail and its respective parts. She is interested in something else: why that particular bobcat is a certain way. And you succeed in explaining it to her, because you point out that practically all bobcats are that way. In so doing, you give a typicality explanation.

Not everyone agrees that typicality facts can explain, however.[3] Resistance to typicality explanation arose, in large part, because there is no rigorous, general account of how typicality facts can be explanatory. If there were such an account, 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.[4] So objections to typicality need not be understood as challenges to the very coherence of the notion; they can be understood as requests for clarification. What is it to be typical? And how do typicality facts explain? Proponents of typicality have also raised questions like these. For example, Goldstein writes that a comprehensive philosophical analysis of typicality explanations 'would be most welcome' ([2012], p. 70).

Hence the present paper, in which I develop a theory of typicality. In Section 2, I analyze typical properties. I also give several mathematical definitions of the notion of 'nearly all' which that analysis invokes. In Section 3, I argue that many explanations invoke facts about properties being typical, I propose a general formalism for typicality explanations, and I say why typicality explanations are explanatory: roughly, they explain by providing concise and informative summaries of the state of the world, by rendering certain phenomena expectable, and by implicitly summarizing facts about what determines what. Finally, in Section 4, I discuss some differences between typicality and probability.

[3]For arguments against typicality explanations of statistical mechanical phenomena, see (Frigg [2011], p. 82) and (Uffink [2007], p. 980).

[4]For example, along with Werndl, Frigg has argued in favor of an interesting topological characterization of typicality ([2012]).

2 An Analysis of Typicality

Pre-theoretically, something is typical just in case nearly all things, of a certain sort, are a certain way. Short-tailed bobcats are typical because nearly all bobcats have short tails. People are typically less than seven feet tall because nearly all people are shorter than seven feet. So there is a close connection between typicality, and nearly all of some things being thus-and-so.

In this section, I make that close connection precise. I propose an analysis of typical properties: an account of what typical properties are. This precisifies the pre-theoretic idea that something is typical just in case nearly all of the relevant things are a certain way.

Here is the 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 bobcats and let P be the property *is short-tailed*. Then P is typical in Γ because nearly all bobcats have short tails.[5]

Typical Property invokes the notion of 'nearly all'. There are many formal definitions of that notion; I cannot cover them all here. So I shall focus on three of the most common.

When Γ is finite, 'nearly all' can be quantified by counting. Let Γ be a large finite set,[6] let P be a property, and let A_P be the set of elements in Γ which exemplify P. 'Nearly all' of the elements in Γ exemplify P $=_{df}$ for some non-negative ϵ much smaller than 1, $\frac{|\Gamma \backslash A_P|}{|\Gamma|} < \epsilon$.[7]

[5] Typical Property can be used to analyze another typicality notion: the notion of a 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 exemplifies P and P is typical in Γ. So a bobcat with a short tail is typical$_o$, relative to the property of being short-tailed and the set of all bobcats.

[6] I stipulate that Γ is large because if Γ were small, it would not make much sense to say that 'nearly all' of the elements of Γ have a certain property. There would not be enough elements in Γ for 'nearly all' of them to be any particular way.

[7] As usual, $|X|$ denotes the cardinality of the set X. The cardinality of a finite set, for example, is just the number of elements it contains. $\Gamma \backslash A_P$ is the set of elements in Γ which are not in A_P.

Call this the 'counting-theoretic' definition of 'nearly all'.

For example, nearly all bobcats have short tails. To see why, let Γ be the set of all (actual) bobcats, let P be the property *is short-tailed*, and let A_P be the set of all (actual) bobcats which exemplify P. For the sake of the example, suppose that the fraction of bobcats which do not have short tails is less than one in a thousand; that is, suppose $\frac{|\Gamma \backslash A_P|}{|\Gamma|} < \frac{1}{1000}$. Since $\frac{1}{1000}$ is quite small, the counting-theoretic definition of 'nearly all' implies that nearly all elements in Γ exemplify P. In other words, nearly all bobcats have short tails.

When Γ is infinite, some other definition of 'nearly all' is required. 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. 'Nearly all' of the elements in Γ exemplify $P =_{df} |\Gamma \backslash A_P| < |\Gamma|$. Call this the 'cardinality-theoretic' definition of 'nearly all'.

For example, on the cardinality-theoretic definition, nearly all real numbers are irrational. Let Γ be the set of reals. Let P be the property *is irrational*, and let A_P be the set of irrationals in Γ. Then the set of elements in Γ which are not in A_P—the set of rationals—is countable. The cardinality of Γ is uncountable. So nearly all reals are irrational.

A third definition of 'nearly all' is used throughout statistical mechanics and quantum mechanics.[8] 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' by measuring size. Let Γ be a set, let P be a property, and let A_P be the set of elements in Γ which exemplify P. Let m be a measure on subsets of Γ: m takes a subset A of Γ as input, and outputs the size $m(A)$ of A.[9] 'Nearly all' of the elements in Γ exemplify P (relative to m) $=_{df}$ for some non-negative ϵ much smaller than 1, $\frac{m(\Gamma \backslash A_P)}{m(\Gamma)} < \epsilon$.[10] Call this the 'measure-theoretic' definition of 'nearly all'.[11]

[8] A version of this definition is discussed in (Frigg [2011], p. 80).

[9] See (Folland [1999]) for the rigorous definition of a measure.

[10] This definition of 'nearly all', like the counting-theoretic definition, is context-dependent: whether or not a particular ϵ counts as 'much smaller than 1' varies from context to context. Consequently, whether or not a property is typical depends on context too. That is a feature of this account of typicality, not a bug: the notion of typicality seems, pretheoretically, like it should depend on context.

[11] This definition requires that A_P be measurable, that $m(\Gamma)$ be non-zero, and that $m(\Gamma)$ be finite.

For example, let Γ be a circle, and let S be a very small sector of Γ; so S is shaped like a very thin slice of pie. Let P be the property of lying outside S, and let A_P be the set of points in Γ which exemplify P. Let m be the standard measure of the areas of two-dimensional shapes. 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 for some very small non-negative ϵ, $\frac{m(\Gamma \setminus A_P)}{m(\Gamma)} < \epsilon$.

Typical Property does not include a free parameter for a measure. So to use the measure-theoretic definition of 'nearly all' in Typical Property, relativize both the analysandum and the analysans 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).

This completes the analysis of typicality. Roughly put, typical properties are exemplified by nearly all members of a given set. Typicality is nearly all.

3 Typicality Explanation

In this section, I argue that facts about typical properties—call them typicality facts—can explain. In Section 3.1, I show that typicality explanations—that is, explanations which cite typicality facts—arise routinely in everyday life. In Section 3.2, I propose a formalism for typicality explanations, and I use my formalism to defend a well-known typicality explanation in statistical mechanics against some objections. In Section 3.3, I say why typicality facts can be explanatory.

3.1 Examples of typicality explanations

Typicality explanations are everywhere. For example, let s be the fact that Mary—the Pennsylvania bobcat which crossed your path—has a tail which is short. Let f be the typicality fact that among bobcats, the property of being short-tailed is typical: so according to Typical Property, f is the fact that nearly all bobcats have short tails. Let b be the fact that Mary is a bobcat. One perfectly good explanation of s cites b and f: Mary has a short tail because she is a bobcat and nearly all bobcats have short tails.

This typicality explanation is distinct from explanations of other types. It is not the sort of explanation which simply cites a cause: the typicality of short-tailed bobcats does not cause Mary to have a short tail. It is not a grounding explanation: the typicality of short-tailed bobcats does not ground the fact that Mary's tail is short. And as I argue in Section 4.1, typicality explanations like these are not probabilistic explanations: neither the explanandum nor the explanans invokes the notion of a short-tailed bobcat being probable. So typicality explanations, though quite commonplace, represent a distinct way of explaining.[12]

Other typicality explanations are just as intuitive as this one. Why does that mosquito not have malaria? Because nearly all mosquitos are malaria-free. Why is that person less than seven feet tall? Because nearly everyone is. Why does that Canadian lynx have such large paws? Because nearly all do.

One might object that in each of these cases, the typicality facts do no explanatory work. In each case, the explanans is really just shorthand for a different explanans that does not invoke typicality facts. Perhaps the typicality fact in the explanation of why Mary's tail is short—the fact that nearly all bobcats have short tails—is shorthand for some causal, biological fact, such as the fact that Mary's genes caused her to have a short tail. So the typicality fact does not explain anything. Rather, it stands for an underlying, non-typicality fact which does all the explaining.

In order to succeed, this objection needs to be spelled out more. Why think that when

[12]Typicality explanations are also distinct from explanations based on generics. 'Mosquitos have malaria' is a generic, but it is not a typicality fact: mosquitos typically do not have malaria. So there is a generic explanation—but not a typicality explanation—of why a particular mosquito has malaria: it is a mosquito, and mosquitos have malaria. 'The property of not having malaria is typical among mosquitos' is a typicality fact, but it is not a generic. So there is a typicality explanation—but not a generic explanation—of why a particular mosquito does not have malaria: it is a mosquito, and the property of not having malaria is typical among mosquitos.

invoked in explanations, typicality facts merely stand proxy for non-typicality facts? Why think that the explanatory power of a putative typicality explanation lies in something other than the facts it explicitly invokes?

But even if these questions can be answered, this objection is problematic for another reason. Set aside the issue of whether the typicality explanation of s which I offered earlier—the one which cites b and f—is genuinely explanatory. Even if not, the typicality explanation offered by the objector—the one which cites Mary's genes—*is plausibly a typicality explanation too*. Unwittingly, the objector rejects one typicality explanation, only to adopt another.

To see why, note that Mary's genes—whatever they are—do not always cause short tails in whatever bobcats have them. Plenty of bobcats with those genes are not short-tailed. Genes can suffer mutation, or be switched off. Tails can be severed by traps or amputated by veterinarians. So at best, Mary has the sort of genes that *typically* cause bobcats to have short tails. Thus, the explanation offered by the objector invokes a typicality fact: the fact that typically, such-and-such genes lead to short-tailed bobcats.

So typicality explanations are not obscure or mysterious. They are intuitively compelling, and we give them all the time. We should, therefore, countenance typicality explanations.

3.2 A formalism for typicality explanation

In this subsection, I propose a formal schema for typicality explanations. Then I discuss a typicality explanation of gas expansion which conforms to the schema, and I defend that explanation against two objections.

The basic schema for typicality explanation is as follows.

$$\begin{array}{rl} & x \text{ is in } \Gamma \\ & P \text{ is typical in } \Gamma \text{ (relative to m)} \\ \hline \therefore & x \text{ has } P \end{array} \tag{1}$$

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 x, Γ, P, m, and the 'nearly all' in the analysis of 'typical' yield different explanations. If that 'nearly all' is not defined measure-theoretically, then the parenthetical in the second line should be dropped. And as for any explanation, in order for an instance of (1) to be explanatory, the premises and the conclusion must all be true.[13]

For example, the explanation of Mary's being short-tailed fits this schema.

$$\begin{array}{rl} & \text{Mary is in the set of all bobcats} \\ & \text{The property } \textit{is short-tailed} \text{ is typical in the set of all bobcats} \\ \hline \therefore & \text{Mary has } \textit{is short-tailed} \end{array} \tag{2}$$

Typicality explanations in many areas of science fit schema (1). As an example, I will discuss an explanation of gas expansion due to Boltzmann ([1877/2015]). Suppose a large box is divided in two by a retractable barrier. A gas occupies the box's left half. Now suppose that the barrier is removed. The gas begins to expand, and relatively quickly, it reaches equilibrium: the gas is evenly distributed throughout the box.

There is a typicality explanation of why the gas reaches equilibrium relatively quickly. It invokes two facts: the gas's microstate X, at the time when the barrier is removed, belongs to a particular macrostate Γ; and the property P of reaching equilibrium relatively quickly is typical in Γ (relative to a particular measure μ). Put roughly, Γ is the initial macrostate of the gas; P is the property of reaching equilibrium within a specified amount of time, by following the dynamics determined by the gas's Hamiltonian; and μ is the modified Lebesgue measure.[14]

So here is the typicality explanation of why the gas reaches equilibrium relatively

[13]Recall that the typicality of a property often depends on context: the counting-theoretic definition of 'nearly all', and the measure-theoretic definition of 'nearly all', are both context-dependent. Because of that, typicality explanations are context-dependent too. And that, I think, is an attractive feature of the proposed theory of typicality. Explanation in general can be context-dependent, so one should expect that typicality explanation can be context-dependent as well.

[14]For rigorous definitions of Γ, P, and μ, see Goldstein ([2001]) and Lazarovici and Reichert ([2015]).

quickly.

$$\begin{array}{rl} & X \text{ is in } \Gamma \\ & P \text{ is typical in } \Gamma \text{ (relative to } \mu) \\ \hline \therefore & X \text{ has } P \end{array} \tag{3}$$

The first premise is true by the definition of Γ. The second premise is true as well, given the measure-theoretic definition of 'nearly all'. The conclusion is true because, as a matter of fact, the gas actually does reach equilibrium relatively quickly.

The literature features several objections to explanations like (3). According to one, such explanations are insufficient because they neglect the Hamiltonian dynamics (Frigg [2011]). But regardless of whether this objection applies to other proffered typicality explanations of gas expansion, it does not apply here. The Hamiltonian dynamics play a crucial role in the rigorous definition of P, and they play a crucial role in selecting μ as the appropriate typicality measure.

According to another objection, such explanations are insufficient because their explanantia do not logically imply their explananda (Uffink [2007]). My response to this objection parallels the response Hempel gave when an analogous objection was made to his view that probability facts can explain. In that case, the objection was that probabilistic explanations are not actually explanatory because in probabilistic explanations, the explanantia do not logically imply the explananda. But as Hempel pointed out, this sort of objection adopts an overly restrictive view of explanation ([1965], p. 391). Many scientific explanations of particular events invoke probabilities. In those cases, the explanantia do not logically imply the explananda, but that is perfectly fine: there are valid probability explanations of particular events. The same sort of response can be given in the case of typicality. Many explanations of particular events rely on the notion of typicality, and thus adhere to schema (1). In those cases, the explanantia do not logically imply the explananda, but that is perfectly fine. As I have shown, typicality explanations are quite common. Logical implication is simply not necessary for explanation.

3.3 An account of explanatory typicality facts

I have offered a formalism for typicality explanations, but I have not yet given an account of what makes typicality facts explanatory. In this section, I argue that typicality facts are explanatory for three related yet distinct reasons: they concisely summarize large amounts of information; they show that certain phenomena are to be expected; and they summarize determination facts—such as facts about what causes what—which are themselves explanatory. But as shall become clear, not all typicality facts seem capable of explaining. So I propose some preliminary criteria for distinguishing explanatory from non-explanatory typicality facts.

To start, consider the first reason why typicality facts explain: they explain because they provide succinct, highly informative summaries of the state of the world. Many typicality facts eliminate lots of possible ways that the world could be, and they do so concisely. For example, let f be the fact that the property *is short-tailed* is typical in the set of all bobcats. Let $\neg f$ be the proposition that it is not the case that *is short-tailed* is typical in the set of all bobcats. There are many more ways for $\neg f$ to obtain than ways for f to obtain;[15] in general, there are far more ways for certain typicality statements to be false than ways for them to be true. So f rules out lots of possible ways that the world could be, and thus, f is a highly informative summary of what the world is like. In addition, f is extremely concise: it expresses all that information in a single, simple sentence.

One might wonder what it takes for a typicality fact to *summarize* the state of the world. I think that typicality facts summarize in three different ways, but for now, let us

[15] As an example, suppose there are only ten bobcats, and suppose that f holds: at most one of the bobcats does not have a short tail, say. There are only 11 ways for that to occur: either none of the bobcats do not have a short tail, or only the first bobcat does not have a short tail, or . . . or only the tenth bobcat does not have a short tail. A simple computation shows that there are 1013 ways for that *not* to occur: the first and second bobcat do not have short tails, the first and third bobcat do not have short tails, and so on. So there are far more ways for $\neg f$ to obtain than for f to obtain.

consider two (I shall discuss the third, which concerns causation, later on). First, some typicality facts summarize by describing what the *actual* world is like. The fact f does this: it summarizes the state of the actual world because it says that nearly all actual-world bobcats have short tails. Second, some typicality facts summarize by describing what nearly all *possible worlds*—in which gases are governed by the same Hamiltonian—are like. As an example, consider the typicality fact invoked in (3): in nearly all possible worlds in which the gas in a box has a specific initial macrostate (and is governed by the same fixed Hamiltonian), that gas reaches equilibrium relatively quickly. This typicality fact summarizes the state of the actual world by placing constraints on what is possible. In other words, this typicality fact summarizes by describing the structure of the set of possible worlds whose gases are governed by a particular Hamiltonian. So the sort of summary provided by this typicality fact is akin to the sort of summary provided by a dynamical law. This typicality fact, like a law of dynamics, describes the actual world by constraining the structure of possibility space.

Because of that, my account of typicality explanation is akin to the deductive-nomological account of explanation. Both accounts hold that certain general rules—typicality facts in my account, laws in the deductive-nomological account—can be explanatory. And both accounts imply that those general rules are explanatory because, in part, of their concision and informativeness.

There are other similarities between my account of typicality explanation and the deductive-nomological account. One concerns the second reason why typicality facts explain: they explain because they render certain phenomena expectable. The deductive-nomological account implies something similar: laws are explanatory, according to the deductive-nomological account, because they show that given the occurrence of various initial conditions, one ought to expect the explanandum (Hempel [1965], p. 337). Likewise for typicality facts: given that x is in Γ, and given that property P is typical in Γ (relative to m), one ought to expect that x has P. Call this a 'typicality norm' of rationality, since it says that typicality facts are a guide to rational expectation.[16]

One might worry about these similarities between the deductive-nomological account of explanation and my account of typicality explanation. In particular, one might worry that my account of typicality explanation is susceptible to the problems facing the deductive-nomological account. For instance, the deductive-nomological account does not always respect the asymmetry of explanation; call this the 'asymmetry problem'. The deductive-nomological account implies, for example, that the length of a flagpole's shadow can be explained by (i) the height of the flagpole, and (ii) a law relating the shadow's length to the flagpole's height. But the deductive-nomological account also implies that the height of the flagpole can be explained by (i) the length of the flagpole's shadow, and (ii) that same law relating the shadow's length to the flagpole's height. And only the first of these seems genuinely explanatory; the second is a deduction, not an explanation. So one might worry that my account of typicality explanation faces some sort of analogous problem.

But it does not. To see why, consider once more the explanation (2) of Mary's short tail: (i) Mary is a bobcat, and (ii) the property *is short-tailed* is typical among bobcats. One might try to get a violation of asymmetry by saying that (ii), along with the fact that Mary is short-tailed, explains the fact that Mary is a bobcat. That is, one might try to get a violation of asymmetry by arguing that the following is an explanation.

$$\begin{array}{rl} & \text{Mary is in the set of all short-tailed creatures} \\ & \text{The property } \textit{is short-tailed} \text{ is typical in the set of all bobcats} \\ \hline \therefore & \text{Mary has } \textit{is a bobcat} \end{array} \tag{4}$$

But this does not have the right structure to be a typicality explanation; it does not conform to schema (1), for instance. Alternatively, one might try to get a violation of asymmetry by saying that Mary's being a bobcat is explained by the following two facts: (i′) Mary has a short tail, and (ii′) the property *is a bobcat* is typical among short-tailed creatures. That is,

[16]This norm only applies if the agent's epistemic state satisfies a variety of other conditions. For instance, it only applies if the agent does not know that x does *not* have P.

one might try to get a violation of asymmetry by arguing that the following is an explanation.

$$\begin{array}{ll} & \text{Mary is in the set of all short-tailed creatures} \\ & \text{The property } \textit{is a bobcat} \text{ is typical in the set of all short-tailed creatures} \\ \hline \therefore & \text{Mary has } \textit{is a bobcat} \end{array} \quad (5)$$

But even though this is an instance of (1), it is still not a typicality explanation. Recall that in order for an instance of (1) to be explanatory, each line must be true. But the second line in (5) is false: bobcats comprise a small fraction of the set of short-tailed creatures.

Now for the third reason why typicality facts explain: they explain because they implicitly summarize causal determination facts—facts, that is, about what causes what—and summaries of causal determination facts are explanatory. For instance, the typicality fact f—that nearly all bobcats have short tails—rules out the many causal histories in which lots of bobcats come to have long tails or no tails at all. Thus, f concisely expresses a large amount of information about the causal structure of the world: that causal structure is such as to make the property *is short-tailed* typical among bobcats. Facts about causal structure are explanatory, and so typicality facts like f—which describe that causal structure—are explanatory too.[17] In short, typicality facts inherit the explanatory capacities of the causal facts they summarize.[18]

One might deny that typicality facts explain by summarizing causal facts. For instance, one might think that the causal summaries encoded in typicality facts are not really explanatory; call this the 'no-explanation' view. Or one might think that typicality facts do not summarize causal facts at all; call this the 'no-summary' view. Either way, on views like

[17]This yields another reason why my account of typicality explanation avoids any problems with asymmetry. Causal facts only explain in one direction: the direction of causation. Typicality facts, insofar as they explain by summarizing causal facts, only explain in one direction too. So even if there is a typicality explanation that—unlike (2)—seems reversible, the 'reverse' candidate explanation will not actually be explanatory, since its direction of purported explanation will go against the direction of causation.

[18]Typicality facts also summarize other kinds of explanatory determination facts. For instance, typicality facts summarize facts about grounding: f, for instance, summarizes facts of the form 'Collections of particles which ground the existence of this particular bobcat also ground the existence of this bobcat's short tail'. Facts about grounding structure are explanatory, so typicality facts like f—which describe that grounding structure—are explanatory too.

these, typicality facts do not inherit the explanatory capacities of facts about what causes what.

I do not accept the no-explanation view, since it contradicts my preferred theories of causal explanation. One such theory, due to Lewis ([1986]), is as follows: to explain an event, it suffices to provide some information about that event's causal history. Typicality facts provide precisely that sort of information: the fact that property P is typical (relative to set Γ) rules out the many causal histories in which a substantial number of elements in Γ do not exemplify P. Another such theory, due to Skow ([2014]), is as follows: to explain an event, it suffices to cite facts about what causal histories would have resulted in the occurrence of some alternative event. More specifically, according to Skow, facts which contain causal-historical information can explain ([2014], p. 455). So again, typicality facts are explanatory.[19]

In addition, I do not accept the no-summary view, because I endorse a particular account of what it takes for a typicality fact to summarize causal facts. To summarize causal facts, according to this account, a typicality fact need only rule out some possible causal histories of the world. In particular, to summarize causal facts, a typicality fact need only imply that at least some causal histories did not obtain.[20]

[19]Some views of causal explanation are incompatible with these accounts. For example, on Cartwright's view, the following candidate explanation is not actually explanatory, even though it summarizes causal-historical information: a particular quail bobs its head when it walks because all quails bob their heads when they walk. This is not explanatory, according to Cartwright, because it misses out on the detailed causal story behind this particular quail's head-bobbing behavior ([1983], pp. 70–1). I agree with Cartwright that the candidate explanation misses out on a more detailed causal story. But I disagree that this is sufficient to render the candidate explanation non-explanatory. The universal generalization 'all quails bob their heads while walking' summarizes causal facts, since it rules out possible causal histories in which quails behave differently. And the detailed causal story behind this particular quails' head-bobbing behavior may well miss out on the more general causal story, partially illuminated by that universal generalization, about why quails behave in that way. Moreover, I see no clear line between explanatory and non-explanatory causal facts. So I am inclined to think that either all causal facts are explanatory, or none of them are. And faced with that choice, I prefer the former: all causal facts can explain, including the universal generalization about quails. If the reader disagrees, then they may consider this an invitation to provide sharp criteria for distinguishing explanatory from non-explanatory causal facts.

[20]The following case reveals a complication in the connection between explanatory typicality facts and causal facts. Suppose a bobcat—call him Avon—loses his short tail in an accident. And suppose that Avon's short tail is surgically reattached. Then consider the following candidate explanation of why Avon has a short tail: Avon is a bobcat, and nearly all bobcats have short tails. One might claim that this is not an explanation: Avon has a short tail because of the surgery, not because nearly all bobcats have short tails. I am

Not all typicality facts seem explanatory, however. Let Γ' be the set containing all short-tailed dogs and also Mary the bobcat. Let f' be the typicality fact that *is short-tailed* is typical in Γ'. One instance of (1) invokes f', rather than f, to explain why Mary has a short tail. The fact f' is informative, renders certain phenomena expectable, and summarizes lots of causal facts. But this instance of (1) does not seem like an explanation, since f' does not seem capable of explaining why Mary's tail, in particular, is short. So what distinguishes explanatory typicality facts like f from non-explanatory typicality facts like f'?[21]

The typicality norm leads to a related problem: by gerryrigging Γ and m, one can use the typicality norm to conclude that one ought to expect contradictory things.[22] For instance, as in (3), let Γ be the set of all microstates belonging to a particular non-equilibrium macrostate, let P be the property of evolving to equilibrium relatively quickly, and let x be a microstate in Γ. Furthermore, let g be the typicality fact that P is typical in Γ (relative to μ). Then by the typicality norm, one ought to expect that x has P. Let Γ' be defined as the set consisting of (i) all microstates in Γ which do not exemplify P, and (ii) the microstate x. And let g' be the typicality fact that $\neg P$ is typical in Γ' (relative to μ), where $\neg P$ is the property of not evolving to equilibrium relatively quickly. Then by the typicality norm, one ought to expect that x has $\neg P$. So the typicality norm implies that one ought to have contradictory expectations.

My solution to both problems relies on naturalness.[23] In order to be explanatory, and

not sure whether this claim is correct. My intuition is that the candidate explanation is indeed explanatory: it is just less explanatory than the explanation which cites the surgery. But if one thinks that the claim is correct, then one owes an account of how the typicality fact in question—that nearly all bobcats have short tails—can explain Mary's short tail but not Avon's. For the sake of brevity, I will not discuss any such account in detail. But here is one that may be worth considering: if a typicality fact explains another fact, then the two facts are grounded in the same kinds of causes. The typicality fact that nearly all bobcats have short tails, and the fact that Avon has a short tail, are not so connected: the typicality fact obtains because of causal facts about the genetics of bobcats, and the latter fact obtains because of causal facts about the surgery (not causal facts about Avon's genetics). So the typicality fact does not explain Avon's short tail because these two facts have different kinds of causal grounds.

[21] An analogous problem—the reference class problem—arises for Hempel's account of probabilistic explanation ([1965], pp. 398–9).

[22] An analogous problem—the problem of ambiguities—arises for Hempel's account of inductive inferences based on probabilities ([1965], pp. 56–7).

[23] Thanks to Audre Brokes for pressing the importance of naturalness here.

in order to guide rational expectation in the way described by the typicality norm, typicality facts must be sufficiently natural. Explanatory, rationality-guiding typicality facts invoke natural sets and natural measures. And the naturalness of a set, as well as the naturalness of a measure, is determined by a variety of factors: for instance, its simplicity, its unity, its homogeneity, and its role in scientific theory.

For example, take the typicality facts f and f'. The set Γ invoked in f—the set of all bobcats—is fairly simple, unified, and homogeneous. It is also invoked in scientific theories, such as evolutionary biology. So f is fairly natural. The set Γ' invoked in f'—the set of all short-tailed dogs and also Mary the bobcat—is fairly complex, disunified, and heterogeneous. And it is not invoked in any scientific theories. So f' is fairly non-natural. Consequently, f can explain and f' cannot.

Similarly, consider the typicality facts g and g'. The set Γ invoked in g—the set of all microstates belonging to a particular non-equilibrium macrostate—is extremely simple, unified, and homogeneous. In addition, Γ plays a crucial role in the formulation of statistical mechanics. So g is extremely natural. But the set Γ' invoked in g'—the set of (i) all microstates in Γ which do not reach equilbrium relatively quickly, and (ii) some specific microstate x—is extremely complex, disunified, and homogeneous. Furthermore, Γ' does not play a crucial role in the formulation of statistical mechanics. So g' is extremely non-natural. Consequently, g can be used to guide rational expectation in accordance with the typicality norm, and g' cannot.

Though I have focused on natural and non-natural sets, the same considerations apply to natural and non-natural measures. In particular, only sufficiently natural measures yield typicality explanations, and only sufficiently natural measures guide rational expectation. The naturalness of the measure μ used to explain why gases expand as they do, for example, is determined principally by the laws. Stationary measures—that is, measures that preserve the sizes of sets over time—are generally more natural than non-stationary measures.[24] That

[24]For a discussion of stationarity, see (Dürr *et al.* [1992]).

is why g is more natural than the typicality fact that P is typical in Γ relative to μ_y, where μ_y is the atomic measure on some specific microstate y which has P. Whereas g invokes a stationary measure, the measure μ_y invoked in the latter typicality fact is non-stationary. So g is the more natural typicality fact.[25]

In short, some typicality facts carve at the joints better than others. Those are the ones on which we focus. Those are the ones we tend to use in everyday conversations. And those are the ones which scientific theories—designed to be simple, and to unify as many phenomena as possible—tend to produce.

4 Differences Between Probability and Typicality

The relationship between probability and typicality is subtle. Their close kinship is articulated by the law of large numbers: according to a particularly plausible interpretation, 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.

So it is worth spelling out the differences between probability and typicality in some detail. As discussed in Section 4.1, there are explanatory differences: some typicality explanations are not probabilistic explanations, and some probabilistic explanations are not typicality explanations. As discussed in Section 4.2, there are formal differences: not all typicality facts can be expressed using probability measures. And as discussed in Section

[25]Stationarity is just one of many factors that can determine the naturalness of a particular measure. On a Humean account of laws, the most natural measure may be whichever best maximizes the simplicity and strength of the system in which it appears. On an anti-Humean account of laws, the most natural measure may be determined by other physical facts. But either way, the situation with typicality measures is akin to the situation with probability measures: just as facts about laws determine the right probability measure for physical theorizing, facts about laws determine the right typicality measure for physical theorizing.

4.3, there are metaphysical differences: to be probable is not always to be typical, and to be typical is not always to be probable.

4.1 Explanatory differences

In this subsection, I discuss the connection between typicality explanations and probabilistic explanations. As shall become clear, these two kinds of explanation have a lot in common. But they are distinct. In particular, there are examples of each kind of explanation which are not examples of the other.

Schema (1) is structurally analogous to Hempel's schema for probabilistic explanation. According to Hempel ([1965], p. 390), probabilistic explanations of particular phenomena take the following form:

$$\begin{array}{rl} & Fi \\ & \underline{p(G,F) = r} \\ \therefore [r] & Gi \end{array} \tag{6}$$

where i is an individual, Fi says that i is F, $p(G, F)$ is the probability that something which is F is also G, and Gi says that i is G. The $[r]$ term in (6) is simply meant to convey that the explanation has a particular strength. And r must be quite high, in order for an instance of (6) to count as an explanation (Hempel [1965], p. 390).

Schemas (1) and (6) are quite similar. The first and third lines of each schema assert facts about particular individuals. The second lines of each schema are somewhat similar too. In (1), the second line says that *nearly all* of the elements in Γ exemplify P (relative to m). In (6), the second line says that *with high probability*, individuals which have F also have G. The similarity derives, of course, from the similarity between 'nearly all' claims and 'has high probability' claims. If nearly all of the elements in Γ exemplify P (relative to m), and one selects an object from Γ at random—and if each object has the same probability of being selected—then there is a high probability that the selected object exemplifies P. Correlatively, if each F has a high probability of being a G—and if there are lots of Fs—then it is extremely likely that nearly all Fs have G.

But typicality explanations and probabilistic explanations come apart. For instance, here is a probabilistic explanation which is not a typicality explanation. Let i be a plutonium atom, let F be the predicate 'is a plutonium atom', and let G be the predicate 'decays'. Let r be the probability that a plutonium atom decays in a fixed amount of time. For the sake of the example, suppose that r is less than $\frac{1}{2}$. And suppose that i actually does decay. Now plug all this into Hempel's schema for probabilistic explanation.

$$\begin{array}{rl} & \text{i is a plutonium atom} \\ & \text{The probability that a plutonium atom decays is r} \\ \hline \therefore [r] & \text{i decays} \end{array} \tag{7}$$

Nowadays, contrary to Hempel, it is generally accepted that some events can be explained by probabilities which are quite low.[26] So (7) counts as a probabilistic explanation.

(7) is not a typicality explanation, however. Since r is less than $\frac{1}{2}$, it is not the case that typical plutonium atoms decay: the majority, in fact, do not. Therefore, the property of decaying is not typical in the set Γ of plutonium atoms. So there is no typicality explanation of the fact that i decays.[27]

On the face of it, many typicality explanations are not probabilistic explanations. Recall the explanation of why Mary has a short tail: she is a bobcat, and nearly all bobcats have short tails. The typicality fact f—that nearly all bobcats have short tails—does not mention probabilities. So this does not appear to be a probabilistic explanation.

One might object that f is indeed a probabilistic fact, though it has been dressed up in the language of typicality. In Section 3.1, I gave several reasons for being wary of objections like these, so I will not rehearse them again. Instead, I will present another

[26]For discussion of the view that low probabilities can explain, see (Strevens [2000], p. 368).
[27]This example is compatible with the view that probabilistic explanations and typicality explanations coincide when the probabilities in question are high. But the following examples challenge that view.

compelling response to this objection. I will present two typicality explanations which, for straightforward reasons, *cannot* be probabilistic explanations.

The first occurs within a version of the Everettian theory of quantum mechanics. According to this version of the theory, 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. Both agree that 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 this 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 a 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 an 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, obey the Born rule.[28] To explain this agreement between observation and the Born rule, suppose we do n experiments on n electrons whose wavefunctions are in a superposition of the electron being on the left and the electron being on the

[28]See (Shankar [1994]) for a detailed discussion of the Born rule.

right.[29] Suppose that n is large. Let Γ be the set of all sequences of observations which we could have made: in one sequence, for example, the electron is always found on the left; in another sequence, the electron alternates between being on the left and being on the right; and so on. Let o be the sequence of observations which we did in fact make: perhaps this sequence goes 'left', 'left', 'right', and so on. Let P be the property, exemplified by some sequences in Γ, of obeying the Born rule: more precisely, sequence x in Γ exemplifies P just in case the relative frequency of 'left' in x is within experimental error of $\frac{1}{2}$.[30] As a matter of fact, o is in Γ and o exemplifies P.

It can be shown that relative to a particularly special typicality measure R, P is typical in Γ. R is special because as Everett demonstrated, R is the only measure which satisfies a series of formal constraints. These constraints arise out of the mathematical structure of pure wave mechanics (Barrett [2017], pp. 33–5), so it is natural to require that the typicality measure satisfy them.[31]

With all that as background, here is a typicality explanation of the fact that the observed outcomes of our experiments conform to the Born rule.

$$\begin{array}{rl} & o \text{ is in } \Gamma \\ & P \text{ is typical in } \Gamma \text{ (relative to } R) \\ \hline \therefore & o \text{ has } P \end{array} \qquad (8)$$

In other words, our observations agree with the Born rule for two reasons. First, the observed sequence of outcomes is one of the many possible sequences of outcomes; this is just the first premise of (8). Second, the property of obeying the Born rule is typical in the set of all possible sequences (relative to R); this is just the second premise of (8).

(8) is a typicality explanation because its second premise is an explanatory typicality

[29] For the sake of the example, suppose that each electron's wavefunction is $\frac{1}{\sqrt{2}}(|L\rangle + |R\rangle)$, where $|L\rangle$ is the state of the electron being on the left and $|R\rangle$ is the state of the electron being on the right.

[30] The Born rule predicts that for an electron in state $\frac{1}{\sqrt{2}}(|L\rangle + |R\rangle)$, the probability of finding that electron on the left is $\frac{1}{2}$.

[31] For more on the role of typicality in Everett's theory, see (Allori *et al.* [2011]).

fact. This typicality fact is explanatory because (i) it is concise and highly informative, (ii) it shows that the explanandum is to be expected, (iii) it summarizes causal facts concerning the unitary evolution of the wavefunction, and (iv) the sets and measures it invokes are sufficiently natural. So this typicality fact can explain.

But (8) is not a probabilistic explanation. For as mentioned earlier, in 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 others ([1956/2012], p. 149). So the sequences of outcomes other than the one invoked in the explanandum of (8)—the sequences other than o—occur too. But in probabilistic explanations, that cannot happen. In probabilistic explanations, *the event invoked in the explanandum is the only outcome, of the various possible mutually exclusive outcomes, that occurs.*

For example, in (7), there are two mutually exclusive events over which the relevant probabilities are defined: the event of i decaying, which has probability r; and the event of i not decaying, which has probability $1 - r$. Since those events are mutually exclusive, they cannot both occur: it is impossible for i to decay and also not decay. So from the fact that (7) is a probabilistic explanation, it follows that each alternative to the decay event described in the explanandum—in particular, the alternative event in which i does not decay—fails to obtain.

This line of thought shows that (8) is not a probabilistic explanation. To see why, let e be the event of getting some specific sequence of observations which contradicts the Born rule: the sequence in which the electron is always found on the left, say. If (8) were a probabilistic explanation, then e and o would be mutually exclusive events; just as the event of i decaying, and the event of i not decaying, are mutually exclusive. But e and o are clearly not mutually exclusive, on this version of Everettian quantum mechanics. For they both obtain. So the explanation provided by (8) cannot be probabilistic.[32]

[32]None of this implies that in Everettian quantum mechanics, probabilistic explanations are impossible. This just shows that (8) is not probabilistic.

Therefore, (8) provides an example of a typicality explanation which cannot be a probabilistic explanation. (8) is a typicality explanation because, put roughly, the typicality fact it invokes is concise, informative, expectation-guiding, connected to various causal facts, and natural. But (8) is not a probabilistic explanation because if it were, then events which contradict the Born rule would not occur. Yet they do.

The second typicality explanation can be extracted from comments made by Elga ([2004]).[33] Let H be the event of a coin landing heads, and let T be the event of that coin landing tails. Consider a universe in which this coin is flipped infinitely many times, producing an infinite sequence $HTTHHTHHHT\ldots$ As Elga points out, it is not clear which probability function best describes this sequence, because according to many different probability functions, the probability of this particular infinite sequence occurring is zero ([2004], pp. 67–9). So which probability function describes this sequence best, and why?

To keep things simple, let us consider just two probability functions. One assigns H a probability of $\frac{1}{2}$ and T a probability of $\frac{1}{2}$; call this function p_1. The other assigns H a probability of $\frac{1}{10}$ and T a probability of $\frac{9}{10}$; call this function p_2. Only p_1, let us suppose, is part of the best deductive system for this world. So p_1, but not p_2, provides a good summary of this world's goings-on. Or as Elga puts it, only p_1 fits the particular sequence of heads and tails which occurs, on a particular definition of 'fits' ([2004], p. 71). Then the question from before is: why? In virtue of what does p_1, rather than p_2, provide a good summary of what happens in this world?

The answer invokes a typicality fact: given a particular measure of typicality endorsed by Elga ([2004], pp. 71–2), according to p_1 but not p_2, the property of satisfying various special conditions—which this particular sequence satisfies—is typical relative to the set of all possible sequences.[34] In other words, p_1 describes this world better than p_2 because p_1,

[33] Though the following explanation does not exactly conform to schema (1), it still uses a typicality fact to explain something. So plausibly, it still qualifies as a typicality explanation.

[34] For example, this particular sequence satisfies the condition of $\frac{1}{2}$ being the limiting relative frequency of H. According to p_1 but not p_2, the property of being an infinite sequence which satisfies that condition is typical.

but not p_2, renders this sequence typical.

This is not a probabilistic explanation. The parallel probabilistic explanation, if there were such a thing, would go like this: p_1 describes this world better than p_2 because p_1, but not p_2, renders this world probable. But this world is not probable, according to p_1. For p_1 assigns this particular sequence a probability of zero.[35]

4.2 Formal differences

There are many 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 invoke the cardinality-theoretic definition of 'nearly all' could be expressed using probability facts concerning m. But it can be shown that this claim is false.[36] 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 invoke the cardinality-theoretic definition of 'nearly all' could be expressed using probability facts. But again, it can be shown that this claim is false.[37]

[35] Thanks to an anonymous reviewer for this point.

[36] The proof is straightforward: given ZFC, there is no function m defined over every powerset of every set Γ.

[37] Again, the proof is straightforward. Let $\Gamma = \mathbb{N}$. For each i, let P_i be the property of being greater than i; so $A_{P_i} = \{i+1, i+2, \ldots\}$. Then for each i, $|\Gamma \backslash A_{P_i}| < |\Gamma|$. If the claim holds, then there is a probability measure m such that for each i, $m(\Gamma \backslash A_{P_i}) < \epsilon$ (for some non-negative ϵ much smaller than 1). But by the upward continuity of measures, $1 = m(\Gamma) = \lim_{i \to \infty} m(\Gamma \backslash A_{P_i}) \leqslant \epsilon$, which is a contradiction.

These technical issues cut to the heart of the formal difference between probability and typicality. The heart of the difference is that while probability measures are always upwards continuous, typicality is often not.[38] It follows that the probability measures of certain typical sets must get arbitrarily small; so even though those sets are typical, they must have arbitrarily low probability.

4.3 Metaphysical differences

The metaphysics of typicality is distinct from the metaphysics of probability. In particular, a thing can be probable without being typical, and a thing can be typical without being probable. So typicality and probability come apart.

To start, here is an example of something which is probable but not typical. Suppose that the probability of a certain type of atom decaying is $\frac{11}{20}$; let Γ be the set of all atoms of this type. Then those atomic decays are probable, insofar as the event of a decay is more probable than the event of a nondecay. But the property of decaying is not typical in Γ; approximately nine in twenty of the Γ atoms do not decay.[39]

Now consider an example of something which is typical but not probable. Imagine a machine that, every second, prints either a 0 or a 1 on a blank square of a tape. Suppose the machine runs for a long, finite amount of time. And suppose that in doing so, the machine produces a sequence which begins with ten thousand 0s, then has a single 1, then has ten thousand more 0s, then a single 1 again, and so on, always repeating the pattern of having ten thousand 0s followed by a 1. Let P be the property of bearing an inscribed 0. Then P is typical, relative to the set of all squares on the tape: nearly all squares contain a 0. But the event of exemplifying P—that is, the event of being a square inscribed with 0—is not probable. For on most accounts of probability, there is a close connection between probability

[38]See (Folland [1999]) for a detailed discussion of upward continuity.
[39]Maudlin makes a similar point ([2011], pp. 316–7).

and randomness: randomness is a core component of what probability is.[40] But the sequence of 0s and 1s is not a random sequence. And so given the connection between randomness and probability, this sequence is not probabilistic.

In other words, while randomness is closely connected to probability, randomness is not closely connected to typicality. Something can be typical without being random. But if something is not random, then it is not probable. So a thing can be typical without being probable.

5 Conclusion

There is much more to say about how the present theory of typicality connects with other areas of philosophical research. The relationship between probability and typicality deserves further exploration: as I intimated in Section 4, the key to that relationship seems to lie in the law of large numbers. The relationship between typicality explanations, and explanations which cite generics, deserves further exploration as well. I also suspect that the notion of typicality can be used to illuminate the nature of special science laws. In addition, typicality seems connected to causation in more ways than I have discussed. And it is worth exploring whether typicality reasoning has its own formal logic.[41]

The present theory of typicality makes way for those future projects. It provides an analysis of typicality in terms of the notion of 'nearly all': typical properties are properties exemplified by nearly all elements of the relevant set. And it provides an account of how typicality facts explain. That analysis of typicality, and that account of typicality explanation, should help facilitate the exploration of the relationship between typicality, probability, generics, special science laws, causation, logic, and perhaps more.

[40] On some accounts of probability—for instance, an actual frequentist account—probabilistic sequences need not be random. And so for those accounts, the event of exemplifying P may indeed be probable.

[41] See (Crane and Wilhelm [forthcoming]) for two formal logics for typicality reasoning.

Acknowledgements

Thanks to David Albert, Valia Allori, Jeffrey Barrett, Audre Brokes, Laura Callahan, Sam Carter, Eddy Chen, Gregory Cherlin, Harry Crane, Roman Frigg, Erikk Geannikis, Verónica Gómez, Mario Hubert, Josh Hunt, Michael Kiessling, Dustin Lazarovici, Joel Lebowitz, Barry Loewer, Tim Maudlin, Paula Reichert, Jonathan Schaffer, Charles Sebens, Michael Strevens, audiences at the 2018 Rutgers Foundations of Probability Seminar and at the Munich Center for Mathematical Philosophy, two anonymous referees, and especially Shelly Goldstein for much helpful feedback and discussion.

References

Allori, V., Goldstein, S., Tumulka, R. and Zanghì, N. [2011]: 'Many Worlds and Schrödinger's First Quantum Theory', *The British Journal for the Philosophy of Science*, **62**, pp. 1–27.

Barrett, J. A. [2017]: 'Typical worlds', *Studies in History and Philosophy of Modern Physics*, **58**, pp. 31–40.

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 and F. Matschinsky, *trans*), *Entropy*, **17**: pp. 1971–2009. (Original work published 1877)

Cartwright, N. [1983]: *How the Laws of Physics Lie*, New York, NY: Oxford University Press.

Crane, H. and Wilhelm, I. [forthcoming]: 'The Logic of Typicality', in V. Allori (*ed.*), *Statistical Mechanics and Scientific Explanation*, Singapore: World Scientific.

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

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

Everett, H. [2012]: 'The Theory of the Universal Wave Function', in J. Barrett and P. Byrne (*eds*), *The Everett Interpretation of Quantum Mechanics*, Princeton, NJ: Princeton University Press. (Original work published 1956)

Folland, G. B. [1999]: *Real Analysis* (2nd *ed*), New York, NY: Wiley & Sons.

Frigg, R. [2011]: 'Why Typicality Does Not Explain the Approach to Equilibrium', in M. Suárez (*ed*), *Probabilities, Causes and Propensities in Physics*, New York, NY: Springer.

Frigg, R. and Werndl, C. [2012]: 'Demystifying Typicality', *Philosophy of Science*, **79**, pp. 917–29.

Goldstein, S. [2001]: 'Boltzmann's Approach to Statistical Mechanics', in J. Bricmont, D. Dürr, M. C. Galavotti, G. Ghirardi, F. Petruccione, and N. Zanghì (*eds*), *Chance in*

Physics, Heidelberg: Springer, pp. 39–54.

Goldstein, S. [2012]: 'Typicality and Notions of Probability in Physics', in Y. Ben-Menahem and M. Hemmo (*eds*), *Probability in Physics*, New York, NY: Springer, pp. 59–71.

Hempel, C. G. [1965]: *Aspects of Scientific Explanation and other Essays in the Philosophy of Science*, New York, NY: The Free Press.

Lazarovici, D. and Reichert, P. [2015]: 'Typicality, Irreversibility and the Status of Macroscopic Laws'. *Erkenntnis*, **80**, pp. 689–716.

Lewis, D. [1986]: 'Causal Explanation', in D. Lewis (ed), *Philosophical Papers, Vol. 2*, New York, NY: Oxford University Press.

Maudlin, T. [2011]: 'Three Roads to Objective Probability'. In C. Beisbart and S. Hartmann (*eds*), *Probabilities in Physics*, New York, NY: Oxford University Press, pp. 293–319.

Shankar, R. [1994]: *Principles of Quantum Mechanics* (2nd *ed*). New York, NY: Springer Science+Business Media.

Skow, B. [2014]: 'Are There Non-Causal Explanations (of Particular Events)?', *British Journal for the Philosophy of Science*, **65**, pp. 445–67.

Strevens, M. [2000]: 'Do Large Probabilities Explain Better?', *Philosophy of Science*, **67**, pp. 366–90.

Uffink, J. [2007]: 'Compendium of the Foundations of Classical Statistical Physics', in J. Butterfield and J. Earman (*eds*), *Philosophy of Physics*, Amsterdam: Elsevier, pp. 923–1074.

Wallace, D. [2012]: *The Emergent Multiverse*, Oxford: Oxford University Press.