A Naturalist's Guide to Objective Chance

Nina Emery*†

I argue that there are such things as nomological probabilities—probabilities that play a certain explanatory role with respect to stable, long-run relative frequencies. Indeed, I argue, we should be willing to accept nomological probabilities even if they turn out to be metaphysically weird or even wholly sui generis entities. I then give an example of one way in which this argument should shape future work on the metaphysics of chance by describing a challenge to a common group of analyses of objective probability—Humean analyses—understood as analyses of nomological probability.

1. Introduction. In the course of everyday reasoning, we often explain the frequency with which some type of event occurs by appealing to the probability of that type of event. Why haven't you ever seen a roulette ball land on the same number twice in a row? Because it is very unlikely for it to do so. Why does an evenly weighted coin land heads roughly half the time? Because the probability of it landing heads is 1/2. Why do you always hit at least one red light on your drive home from work? Because the probability of hitting at least one red light is so high.

You might think that this sort of reasoning should not be taken too seriously. You might think that using probabilities to explain the frequencies of everyday macro-physical events is a mistake or at most a useful shorthand, a placeholder of sorts for some more complicated nonprobabilistic explanation. But a similar sort of reasoning also plays a crucial role in some of our best scientific theories. And insofar as it does, it is not so easily dismissed.

Received July 2016; revised October 2016.

* To contact the author, please write to: Brown University, Department of Philosophy; e-mail: nina_emery@brown.edu.

† Many thanks to Carl Hoefer, Ned Hall, Brad Skow, Josh Schechter, Chris Meacham, Jennan Ismael, Elizabeth Barnes, Carla Merino, Amy Seymour, Jenn Wang, Mauricio Saurez, Al Wilson, Matt Duncan, and audiences at the Young Philosophers Lecture Series at DePauw University and the 2015 Society for the Metaphysics of Science Conference.

Philosophy of Science, 84 (July 2017) pp. 480–499. 0031-8248/2017/8403-0005$10.00

I will use the name 'nomological probability' for the type of probability that plays the explanatory role just described.

> **Definition of nomological probability.** If (P) is a true proposition of the form *the probability of an event of type E occurring in a situation of type S is x*, and (F) is a true proposition of the form *the stable, long-run relative frequency with which events of type E occur in situations of type S is y*, and *x* is close to *y*, then (P) explains (F) if and only if the probability referred to in (P) is a nomological probability.[1]

'Nomological' probabilities are so named because the explanatory role that such probabilities play with respect to stable, long-run relative frequencies is reminiscent of the explanatory role that laws play with respect to associated regularities. If it is a law that events of type E never occur, that explains why events of type E never occur. If it is a law that all Fs are Gs, that explains why all Fs are Gs. Nomological probabilities, then, are probabilities that play a lawlike role.

In this paper, I argue that we ought to believe that there are such things as nomological probabilities even if the further metaphysics we are able to give of such probabilities turns out to be problematic in various ways. In particular, I will argue that we ought to believe that there are nomological probabilities even if it turns out that the only successful analysis of such probabilities involves entities that are supposed to be in some sense metaphysically weird, or even if there is no successful analysis of such entities at all, and nomological probabilities turn out to be sui generis. I then give an example of one way in which this argument should shape future work on the metaphysics of chance. In particular, I describe a challenge associated with understanding one common group of analyses of objective probability—Humean analyses—as analyses of nomological probability.

Of course, the idea that some sort of objective probability plays an important role in our best scientific theories is hardly new. But the particular explanatory role that I will identify has seen relatively little discussion in the literature,[2] and its implications have largely gone unrecognized. Insofar as

1. Note that this definition is wholly silent on whether (P) also explains other features of the world, like the fact that the collection of events of type E displays a certain sort of randomness. It is also wholly silent on whether there is an explanatory relation between (P) and (F) if *x* is not close to *y* and, if there is not, how close *x* needs to be to *y* before the explanatory relation obtains. I say more about this in n. 14. Note that here and in what follows I assume that probabilities attach to events, as opposed to propositions. Nothing important hangs on this and the relevant event-talk can be easily reinterpreted in terms of propositions if the reader so desires.

2. An exception is Hájek (1996), who uses the claim that probabilities must explain frequencies as a premise in an argument against actual frequentist views of probability. Pro-

philosophers acknowledge that we want objective probabilities to explain frequencies, they seem to think that desideratum is just one among many that compete to influence our metaphysics of chance—one that might be trumped, for instance, by concerns about the type of entities that explanatorily powerful probabilities would be and how those entities fit into our broader metaphysics. But given the argument I present below, this is a mistake. We should think that there are nomological probabilities even if we are forced to accept that such probabilities violate various principles we had hoped our metaphysics would respect and even if such probabilities do not fit into our broader metaphysics at all. Insofar as the existence of such probabilities raises a challenge for Humean analyses of objective probability, for instance, it is not a challenge that can be postponed or dismissed on the grounds that the Humean approach's virtues outweigh any weakness that the challenge exposes.

Two clarifications before we begin. First, as the title says, this is supposed to be a *naturalist's* guide to objective probability. As I will be using the term here, to be a naturalist is to take standard scientific methodology as a good guide to successful inquiry into what the world is like. Insofar as one is a naturalist, then, one should think that the norms and constraints governing scientific theory choice should govern metaphysical theory choice as well. This does not mean, however, that one must take everything that scientists claim at face value. Nor does it mean that one should think that philosophers have nothing to add to the understanding of scientific practice. Naturalism is one thing. Scientism is another.

Second, as the title of the article says, this is supposed to be a naturalist's guide to *objective* chance. I take it that it is obvious, even just from what has been said so far, that whatever else they are, nomological probabilities are objective probabilities—probabilities that are determined by physical features of the world independently of any particular epistemic position that we might occupy within that world. However much we might wish it to be otherwise, the behavior of roulette wheels and coin tosses is not explained by facts about the beliefs or evidence we happen to have. You might, of course, explain our expectations regarding how roulette wheels and the like behave by appealing to facts about the beliefs or evidence that we have, but

pensity theorists have also sometimes claimed it to be an advantage of their theories that probabilities are explanatorily powerful (see, e.g., Miller 1995; Gillies 2000; as an example of prominent propensity theorist who makes no mention of explanatory power, consider Popper [1959]). The explanatory power of chances was also clearly an important influence on Humphreys (see Humphreys 1985) and on recent work by Suárez (see especially Suárez 2013)—though note that both of these authors distinguish propensities (which are explanatorily powerful) from merely descriptive probabilities. The explanatory role of probabilities is also frequently mentioned in discussions of deterministic chance. See, e.g., Albert (2000), Loewer (2001), Meacham (2010a, 2010b), and Emery (2015).

explaining our expectations regarding some phenomena and explaining the phenomena itself are two very different things.[3]

So the argument presented in the first part (secs. 2 and 3)—the argument for the existence of nomological probabilities—is also an argument for the existence of objective probabilities, or chances. (I use the terms 'objective probability', 'chance', and 'objective chance' interchangeably.) But nowhere in what follows do I commit myself to the view that nomological probability is the only type of objective probability. Insofar as you are a naturalist, the arguments below should convince you to believe in nomological probabilities. If, in addition, you have antecedent metaphysical commitments or intuitions that commit you to some notion of objective probability that comes apart from nomological probability, by all means feel free to accept such objective non-nomological probabilities into your metaphysics. But you must accept them in addition to nomological probabilities. Or so I will argue.

2. The Argument for Nomological Probability. Why think that there are such things as nomological probabilities? In short, because they play a crucial explanatory role in our best scientific theories. In particular, they explain certain patterns of events that would otherwise go unexplained.

Here is a paradigm case of the sort I have in mind. Consider carbon-11, a radioactive isotope of carbon that decays into boron. According to the standard interpretation of quantum mechanical phenomena—the interpretation taught in physics textbooks and in physics classrooms—the fundamental laws governing the decay of carbon-11 atoms are indeterministic. More specifically, there is no feature F such that all and only the carbon-11 atoms that have F will decay within a given time period. Nonetheless, over a long series of trials, we observe a relatively stable pattern in the decay of such atoms. And we take this relatively stable pattern in our data to be indicative of an even more widespread pattern in the sorts of events that occur in similar circumstances, namely (1):

> (1) For the vast majority of samples of carbon-11 atoms, around half of those atoms decay within 20 minutes.[4]

3. It is worth emphasizing that objective probabilities—as I am using the term—are something over and above rational degrees of belief. The degrees of belief that it is rational to have in a given situation may depend on objective features of the world, but they also depend on features of your epistemic position, like what other beliefs you hold, what evidence you happen to have, and what sort of being you are. These are paradigmatically subjective features of the world. Note that there are some philosophers who draw the objective/subjective probability distinction somewhat differently. As an example, Williamson (2010) uses 'objective probability' to refer to degrees of rational belief.

4. Here we should understand the relevant sample of carbon-11 atoms as being sufficiently large. Insofar as you consider smaller samples, it will still be the case that most

The standard explanation for (1)—the explanation that appears in physics textbooks and that is taught in a physics classrooms—is a probabilistic explanation. It is standard, in other words, to explain (1) by appeal to (2):

> (2) For any particular carbon-11 atom, the probability of that atom decaying within any particular 20-minute interval is 1/2.

One way to argue for the existence of nomological probabilities, then, would be to take this standard explanation at face value. Scientists assert that (2) explains (1). So (2) explains (1). So the probability referred to in (2) is a nomological probability.

Insofar as you find that argument convincing, great! You are already on board with the existence of nomological probability. But insofar as you are not convinced, there is another way of arguing for the existence of nomological probability, one that takes its cue not just from what scientists say about particular cases like the one described above, but from a more general feature of standard scientific practice. That argument goes like this: assuming that the standard interpretation of quantum mechanical phenomena is correct, if (2) does not explain (1), then nothing does. And (1) describes the sort of pattern such that leaving that pattern unexplained is a serious cost. In particular, the costs associated with leaving (1) unexplained are higher than the costs associated with using (2) to explain (1). So, we ought to adopt a theory according to which (2) explains (1). So, the probability referred to in (1) is a nomological probability. This is the argument that I will focus on here. I will refer to it as *the* argument for nomological probability.

Why think that if (2) does not explain (1), then nothing does? Because the only other other candidate explanation for (1) is:

> (3) There is some feature F, which roughly half of all carbon-11 atoms have, and which is such that all and only carbon-11 atoms that have F decay within a particular 20-minute interval.

And, as I mentioned above, according to the standard interpretation of quantum mechanical phenomena, (3) is false.

Why think that (1) describes the sort of pattern such that leaving that pattern unexplained would be a serious cost? The first step in answering this question is to see that (1) describes what I will call a *robust pattern* of events—a pattern that holds under a variety of temporal, spatial, and counterfactual conditions. It does not matter whether you consider a sample of carbon-11 before 1950, or between 1950 and 1975, or after 2000, or whether

samples exhibit the relevant pattern, but it may not be the case that all, or nearly all of them, do.

that sample was observed in Austria or Australia, or whether it contained a few additional atoms or a few fewer. The pattern described in (1) is supposed to hold under all of these—and many more—conditions.[5]

Why think that it is a serious cost to leave robust patterns, like (1), unexplained? Because it is a widespread and well-established part of scientific practice to view robust patterns in this way. Scientists do not just accept that certain kinds of collisions between particles are reliably followed by certain patterns in the data collected at the Large Hadron Collider—they adopt a theory according to which those patterns are explained by short-lived and otherwise unobservable particles that result from those collisions. They do not just make note of slight periodic oscillations in the light emitted from distant stars—they adopt a theory according to which those oscillations are explained by a nearby exoplanet. None of our best scientific theories, whether at the level of fundamental physics or in the special sciences, leave robust patterns unexplained. Why not? Presumably, because to do so would be considered a serious cost.[6]

But it is not enough to establish that (1) describes the sort of pattern such that leaving that pattern unexplained would be a serious cost. We also need to show that the costs associated with leaving (1) unexplained are higher than the costs associated with using (2) to explain (1). The costs associated with using (2) to explain (1) depend on the metaphysical account that one gives of nomological probabilities. I will say more about various candidate accounts in the second part of the paper (sec. 4). For now it will suffice to note that a metaphysical account of nomological probabilities could be costly in two very broad sorts of ways.

First, it could be metaphysically weird. Nomological probabilities could turn out to be the sorts of entities such that, even if they play a role in folk

5. Note that I take the pattern to be explained—the pattern described by (1)—to be one that goes beyond the pattern that is observed in the data. I do so because I think this is more representative of actual scientific practice. Scientists take the pattern observed in the data to be evidence of a more comprehensive pattern, one that holds under a variety of different temporal, spatial, and counterfactual conditions, and look for a way to explain that more comprehensive pattern. Insofar as you think that scientists are instead primarily concerned with explaining patterns only in actual phenomena, you are welcome to reinterpret the discussion in that way. In that case the patterns to be explained will be temporally and spatially but not counterfactually robust. It will be important for the arguments that follow that one not adopt the view that only patterns in the actually collected data require explanation. But I take it that that view is anyway implausible since it makes whether a pattern requires explanation depend on highly accidental features of when, where, and how we happen to construct our experiments. I say more about this in the discussion of actual frequentism in sec. 4.

6. How robust does a pattern need to be in order to be such that leaving it unexplained is a serious cost? That is a good question, but one that I will set aside here. My claim is just that the sort of pattern described by (1) is a clear case of such a pattern.

theory or ordinary language, we should strive to eliminate from our metaphysics wherever we can. What counts as a metaphysically weird entity will, of course, depend on your background metaphysical commitments. But the obvious way in which nomological probabilities could turn out to be metaphysically weird by the lights of many contemporary metaphysicians is that they could turn out to be non-Humean—they could involve primitive modal, causal, or dispositional properties. According to Humeans, at least, such properties are metaphysically suspect; insofar as we can do without them, we should.[7]

But to be metaphysically weird is not the only way in which an entity can be costly. Entities can also be costly because they are novel—because they cannot be wholly or partly identified with or analyzed in terms of other, more familiar entities. Novel entities, most philosophers agree, are a significant cost; insofar as we can do without them, we should.

Novelty, of course, comes in degrees, and the extent to which an entity is novel may be controversial. But by way of illustrating the difference between metaphysical weirdness and metaphysical novelty, consider various potential propensity analyses of nomological probabilities, according to which nomological probabilities are some sort of disposition or tendency for a certain type of physical system to produce a certain type of result.[8] These accounts are non-Humean, and thus, at least by the Humeans' lights, metaphysically weird. But they do not incur any especially significant further costs associated with being novel. Dispositions and tendencies are the sorts of things that are familiar from everyday language, from folk theorizing, and from other areas of philosophical inquiry. So even if the propensity interpretation does not provide us with a fully reductive or Humean analysis of nomological probability, it identifies probabilities with a broad, relatively familiar class of metaphysical entities. On a propensity account, then, nomological probabilities might count as weird, but they would not be especially novel.

The sort of analysis that would incur significant costs associated with novelty is an analysis on which nomological probabilities are sui generis entities, which cannot be wholly or partly identified with or analyzed in terms

7. Throughout what follows I will use non-Humean entities as my go-to example of entities that might be thought of as metaphysically weird. But note that many non-Humeans will not find them weird at all, and that I am not endorsing or arguing for the view that non-Humean entities are in fact weird. As an example of other types of entities that might be thought of as metaphysically weird, consider unobservable entities, which would be weird by the lights of certain kinds of empiricists, or mental entities, which would be weird by the lights of certain physicalists.

8. For details, see Popper (1959) and Gillies (2000). I assume here and throughout that propensity theorists do not go on to give a Humean account of dispositions or tendencies.

of other, more familiar classes of entities.[9] On a sui generis account of nomological probabilities, the explanatory role that nomological probabilities play would be our main and perhaps only way of grasping what such entities are. We cannot say what nomological probabilities are. We can only say what they do.

All that is by way of spelling out two broad ways in which nomological probabilities might turn out to be costly. The crucial point here is that even if nomological probabilities turn out to be costly in one or both of these ways, we still ought to think that they exist. Let's focus first on the possibility that the only successful analysis of nomological probabilities involves entities that are supposed by some to be metaphysically weird. It is clearly a widespread and well-established part of scientific practice not to leave robust patterns—like the one described in (1)—unexplained even if the only plausible explanation for them requires introducing metaphysically weird entities. The history of scientific inquiry is replete with examples of scientists allowing into their metaphysics entities that would previously have been thought of as strange or suspect, as long as those entities play a robust explanatory role. Think of Newton positing a force that operates on heavenly and earthly bodies alike. Or Planck suggesting that light exhibits properties associated with both a particle and a wave. Or contemporary physicists positing fundamental particle after fundamental particle in order to explain their experimental results, despite the fact that their standard model is becoming increasingly unwieldy.[10]

In each of these cases the need to explain a robust pattern in the phenomena trumps concerns about metaphysical weirdness. And with examples like these in mind, it is hard to imagine how the worry that nomological probabilities may not admit of, for instance, a Humean analysis, could justify the

9. One example of an approach to objective probability that might plausibly be understood as falling within this category is the one found in the first section of Maudlin (2007b). Note, however, that the later sections of that paper describe theories of deterministic objective probability that do appear to admit of at least partial analysis in terms of more familiar entities and thus do not count as novel, as I am using that term. As an example of someone who puts forward a view on which deterministic, macro-level probabilities do not admit of any further analysis, consider Sober (2010).

10. As a referee has helpfully pointed out, these examples put some pressure on the distinction between mere metaphysical weirdness and novelty. Does Planck's theory of light, for instance, involve a merely weird entity or a novel entity? The key thing to note is that, as I have defined these terms above, novelty implies weirdness. Metaphysically weird entities are entities of which we should be suspicious, and I take it that everyone agrees that we should be suspicious of novel entities. In any case, insofar as you think we should be willing to introduce genuinely novel entities, you presumably also think that we should be willing to introduce merely metaphysically suspicious entities. So you should accept the conclusion of my argument even if you take these cases to involve genuine novelty as opposed to mere weirdness.

attitude that such entities are just too strange to be allowed into our metaphysics. If we use scientific practice as our guide, and if the alternative is leaving a robust pattern unexplained, it seems that we ought to swallow our metaphysical scruples and introduce whatever entities necessary—whether Humean or not.

This may sound surprising. Philosophers who think of themselves in the empiricist tradition, and who place significant constraints on the kinds of entities they allow into their metaphysics, also tend to think of themselves as being more scientifically respectable than their opponents. What the argument above shows is that, in at least some cases, by invoking the very constraints that are supposed to make them scientifically respectable, such philosophers risk running afoul of actual scientific practice. According to Lewis (1994, 474), the motivation for Humeanism is the desire to, "resist philosophical arguments that there are more things in heaven and earth than physics has dreamt of." But surely those who think that we should resist philosophical arguments that there are more things in heaven and earth than appear in our best physics should also think that we should resist philosophical arguments that there are fewer things in heaven and earth than appear in our best physics. As Maudlin (2007a, 77) puts it, "Philosophical accounts that force upon us something that physics rejects ought to be viewed with suspicion. But equally suspect are philosophical scruples that rule out what physics happily acknowledges." And, as the above argument shows, Humeanism runs the risk of turning on just such scruples.

So the argument for the existence of nomological probabilities goes through even if the only successful analysis of nomological probabilities turns out to be non-Humean (or otherwise (supposedly) metaphysically weird). The same is true even if the only plausible account of nomological probabilities turns out to be one on which such probabilities are also novel. After all, scientists not only allow metaphysically strange entities into their theories as long as those entities explain robust patterns that would otherwise go unexplained—they also allow sui generis entities in order to play the same role.

Here is an example.[11] Astronomers have observed that the rate of expansion of the universe is accelerating, and physicists have ruled out all possible explanations for this acceleration that are based on those entities to which our best scientific theories are already committed. Given that no other explanation is in the offing, physicists have, by and large, reluctantly decided to posit a new sort of entity—dark energy. What is dark energy? Physicists can-

11. Other examples that could be used to make a similar point include the electromagnetic field as it was introduced in the nineteenth century, and the wave function as it plays a role in standard approaches to nonrelativistic quantum mechanics. See Maudlin (2013) for a discussion of the former and how it impacts the way that we should think about the latter.

not tell us. They can only tell us what dark energy does. It causes the rate of expansion of the universe to accelerate.

Now it may be that in the future we will learn more about what dark energy is—we will be able to do more than identify its explanatory role. Indeed it may be that one of the key desiderata of future theories is that they be able to provide further insight along such lines. But the inclusion of dark energy in our best scientific theories does not appear to be conditional on that kind of future success. Scientists are, by and large, confident that dark energy exists—if it does not, there is no explanation for the rate at which the expansion of the universe is accelerating, and that is the sort of fact for which there must be some sort of explanation. If it turns out that we need to accept dark energy as a sui generis entity, so be it.

My claim is that philosophers should be willing to take the same attitude toward nomological probabilities. Given the crucial explanatory role that they play, we should be more convinced that nomological probabilities exist than of any particular story about what they are. In fact, we should be open to the idea that there is no further story about what they are. It may turn out that we cannot say what nomological probabilities are; we can only say what they do. Nomological probabilities are those objective probabilities that explain associated frequencies in the way described above.

Let's recap. I have argued that on the standard interpretation of quantum mechanical phenomena, either (2) explains (1) or nothing does. And I have argued that the following is an important principle governing standard scientific practice.

> **The pattern-explanation constraint.** Insofar as the only way to avoid leaving a robust pattern unexplained is to introduce a type of entity that is metaphysically weird or novel, we ought to introduce such entities.

Insofar as we are naturalists (in the sense described in the introduction), it follows that we ought to use the pattern-explanation constraint as a guide to metaphysical theory choice as well. So even if the only plausible analysis of nomological probabilities involves entities that are supposed to be in some sense metaphysically weird—even if, for instance, the only plausible analysis is non-Humean—we still ought to accept that there are nomological probabilities. Indeed, even if the only plausible account of nomological probabilities is one on which they are sui generis, we ought to accept that there are nomological probabilities.[12]

12. It is worth emphasizing that nothing that I have said above establishes that nomological probabilities will have to be in any sense weird or novel. My point here is just that even if they are, we should still think that there are such things.

In section 4, I will demonstrate one way in which this argument should shape future work on the metaphysics of chance. But before I do, let me briefly address three objections that might be thought to create difficulty for the argument for nomological probability as I have presented it above.

3. Three Objections. The first two objections I will consider are to do with my appeal to the standard interpretation of quantum mechanical phenomena in the argument for nomological probability. That argument relied on the premise that if (2) does not explain (1), then nothing does. The reason I gave in support of that premise was that the only other plausible explanation of (1) was (3), and according to our best scientific theories—in particular, according to the standard interpretation of quantum mechanical phenomena—(3) is false.

Consider first the fact that the standard interpretation of quantum mechanical phenomena—the one taught in physics textbooks—is not obviously one of our best scientific theories. In particular, the dynamics posited by that standard interpretation is deeply problematic. It relies on a fundamental distinction between measurement processes, which cause a certain type of evolution, called *collapse*, and nonmeasurement processes, which do not. Doesn't this undermine the argument for the existence of nomological probabilities?

In fact, it does not. The argument for nomological probabilities goes through even if one eschews the standard interpretation and adopts instead some version of the Ghirardi-Rimini-Weber (GRW) interpretation, according to which there is no fundamental distinction between measurement and nonmeasurement process; instead collapses are the result of a spontaneous, indeterministic process. On a GRW interpretation, the fundamental laws are still indeterministic and (3) is still false. And versions of the GRW interpretation are clearly among the live candidates for our best interpretation of quantum mechanical phenomena.

This gives rise to a related objection, however. Although versions of the GRW interpretation are clearly among the live candidates for our best interpretation of quantum mechanical phenomena, they are not the only live candidates. Bohmian mechanics is also a live candidate, and according to Bohmian mechanics, (3) is true.[13]

How does the argument for the existence of nomological probability fare once we acknowledge that Bohmian mechanics is also a live candidate for our best scientific theory of quantum mechanical phenomena? It still estab-

13. Few, if any, believe that Bohmian mechanics is the only live interpretation. Insofar as there is any consensus, it is that Bohmian mechanics is a less plausible interpretation than the GRW theory because there is no relativistic extension of Bohmian mechanics, whereas there is a relativistic extension of some versions of GRW.

lishes a substantive conclusion—it establishes that there are nomological probabilities in the region of logical space that is currently compatible with our best scientific theories. And it still follows from the argument that our metaphysics needs to leave room for the existence of nomological probabilities—we ought not take on commitments from which it would follow that there are no such things, or on which such things would be especially difficult to countenance or understand. To do so would be to settle, by stipulation, what is otherwise an open scientific question. So the existence of Bohmian mechanics as an alternative interpretation does not undermine the importance of the argument for nomological probability.

Those were the first two objections. The third objection that I will consider is of a more general philosophical nature. So far I have said nothing whatsoever about a theory of explanation. But surely we cannot establish that (2) explains (1) without first establishing under what conditions one thing explains another.

In fact, there are good reasons to think that we can do just that. For one thing, since satisfying the pattern-explanation constraint warrants the introduction of weird (or even novel) entities in order to explain robust patterns, positing that (2) explains (1) places only minimal constraints on one's theory of explanation. If there is no way of plausibly interpreting (2) as being both explanatorily robust and about familiar entities, one can simply reinterpret (2) as in fact being about some novel entities instead. But note that it is also plausible that the considerations above support not only the introduction of novel entities, but also the introduction of novel explanatory relations. Insofar as our favored theory of explanation gives us no reason for thinking that (2) explains (1) on any interpretation of (2), then we might simply conclude that there exists a novel sort of explanatory relation, which was previously unrecognized in our theorizing, and which holds between (2) and (1). Such a novel explanatory relation would surely be a cost.[14] But it appears to be precisely the sort of cost that is trumped by the pattern-explanation constraint. We do not know what dark energy is, so we cannot be sure that the way in which it explains the accelerating rate of expansion of the uni-

14. How much of a cost this novel explanatory relation would be will depend on how complicated the explanatory connection is. Consider, for instance, the use of 'close to' in the definition of nomological probability. Is there some absolute value that this refers to? Or some value that varies depending on, for instance, how many times the relevant set-up has produced a relevant outcome? (The latter would need to be true insofar as the instances in which (P) explains (F) are going to track those instances in which (F) confirms (P).) Is there also an explanatory relationship between (P) and (F) in cases where the relative frequency is not at all close to the chance? These are all features of the explanatory role of chance that have been left unspecified. And they may add to the extent to which stipulating that (2) explains (1)—insofar as we in fact need to do so—is a cost. Many thanks to Carl Hoefer for pushing me on this point.

verse is compatible with any particular theory of explanation. But that does not seem to worry the scientists who posit dark energy. They are more convinced of the fact that something must explain the accelerating rate of expansion of the universe than they are of any particular story about what explanation is. I suggest that as metaphysicians we take the same attitude toward the explanatory role of nomological probability. We should be more convinced of the fact that something explains (1) than we are of any particular story about what explanation is.

4. What Are Nomological Probabilities? In this section, I will illustrate one way in which the argument in section 2 should shape future work on the metaphysics of chance. In particular, I will present a challenge to one group of analyses of objective probability—Humean analyses—understood as analyses of nomological probability.

Humeans think that all that exists at the fundamental level is the Humean mosaic—the distribution of categorical properties across space and time.[15] Insofar as there are facts about what caused what, or what could have happened, or what was disposed to occur under various circumstances, those facts depend, in some important sense, on the facts about the mosaic.[16]

Insofar as Humeans think that everything depends on the mosaic, they must think that the nomological probabilities depend on the mosaic. As a result, Humeans face an immediate worry. In broad strokes, the worry is this: Humeans think that the nomological probabilities depend on the mosaic. But they must—given the argument in section 2—also think that the nomological probabilities explain features of that mosaic. And how can it be that nomological probabilities explain features of the mosaic if the nomological probabilities themselves depend on the mosaic?

Call this *the explanatory challenge*. Surprisingly, although versions of this challenge are familiar from the debate over Humean analyses of laws of nature, whether and how these challenges extend to the debate over Humean analyses of chance has received relatively little attention.[17] Different versions of the explanatory challenge will target different Humean analyses of objective probability. The most straightforward Humean analysis is an actual frequency analysis, according to which the objective probability of

15. See the introduction to Lewis (1986); and also see Lewis (1994).

16. The precise sense of dependence at issue here will vary depending on what version of Humeanism you have in mind. According to Lewis, for instance, the latter supervene on the former. But a contemporary Humean might think, for instance, that the causal, modal, and dispositional facts are grounded in the mosaic. I am using the term 'depends' here in a way that is supposed to be neutral between these options.

17. On the explanatory challenge against Humean analyses of laws, see, e.g., Armstrong (1985) and Maudlin (2007a).

an event of type E happening in a situation of type S is just the relative frequency with which events of type E actually occur in situations of type S.

Understood as an analysis of nomological probability, the actual frequency analysis faces a particularly serious version of the explanatory challenge. After all, according to that sort of analysis, (2) is equivalent to (1). So (2) cannot explain (1). Nothing can explain itself.

The actual frequentist may be tempted to claim that even if (2) does not explain (1) it explains parts of the pattern described in (1). For instance, it explains (4):

> (4) For any particular sample of carbon-11 atoms that we have observed, around half of those atoms decay within 20 minutes.

But this is a non sequitur. If the discussion in section 2 is correct then it does not matter whether the actual frequentist can show that (2) explains (4). Given the argument in section 2, they also need to show that (2) explains (1). Because if (2) does not explain (1) then nothing does, and (1) describes the sort of pattern that requires an explanation.

Of course, the actual frequentist might suggest that the argument in section 2 was mistaken. It is not patterns like the one described in (1) that require explanation at all. What require explanation are patterns like the one described in (4). But surely that cannot be correct. To say that (4) requires explanation while (1) does not is to say that whether a pattern requires explanation depends on whether we happen to have observed all instances of that pattern. But whether we happen to have observed every instance of some particular pattern is highly accidental. It depends on where we happen to have constructed our experiments, when we happen to have run them, whether we happen to have been paying attention on any particular occasion, and so on. And surely whether some otherwise robust pattern requires explanation cannot be so accidental.[18]

That the actual frequency analysis faces this sort of explanatory challenge has been noted before (see Hájek 1996). What is important to observe here is the strength of the challenge. It looks like one cannot give an actual frequency

18. Note that it is perfectly compatible with what I say here to think that (4) requires explanation. What is important is just that you not think that only patterns like (4)—i.e., patterns that are actually observed—require explanation. Perhaps a more plausible claim is that the difference between (4) and (1) is not that (4) is about a pattern that we have observed whereas (1) is not, but rather that (1) is about a complete pattern (one that includes all actual instances of a type of set-up), whereas (4) is not. But whether or not a pattern is complete in this sense also seems highly accidental. It depends, for instance, on whether there was one more instance of the relevant set-up in some distant region of spacetime. Surely whether some robust pattern requires explanation cannot be so accidental.

analysis of nomological probability. Where does that leave those who are otherwise inclined toward such analyses? Should they give up on the existence of nomological probability? They should not. Should they attempt to engage the rest of us by listing the other virtues of their analysis or the costs of the relevant alternatives? Again, no. The argument in section 2 shows that we should think that there are nomological probabilities, pretty much regardless of the costs associated with accepting such entities into our metaphysics. Insofar as you are a committed actual frequentist, then, you cannot simply dismiss the explanatory challenge outlined above. You must meet it. Until you have done so, your position is untenable.

Of course, many philosophers already think actual frequentism is untenable. The explanatory challenge becomes more interesting, therefore, when it targets other, more prominent Humean analyses of objective probability.[19] Perhaps the most prominent such analysis is the best systems analysis or the BSA.[20] The BSA is first and foremost an analysis of laws. It says that (i) the best system is the set of propositions describing the Humean mosaic that best balances simplicity and strength and (ii) the laws are those propositions that appear in the best system.

19. A Humean analysis that I do not have space to discuss in detail here is Hoefer's (2007) Humean objective chance. According to Hoefer, "Chances are constituted by the existence of patterns in the mosaic of events in the world" (564). But while some chances are "simply there, to be discerned, in the patterns of events," other chances are "regularities that are guaranteed by the structure of the . . . chance setup." In the former case, chances just amount to patterns of events. In the latter case, they are patterns, but patterns that are the result of some underlying features of the system in question. Here, in brief, is why I do not think that Hoefer's account works as an analysis of nomological probability. Call the first group described above the basic Humean chances and the second group the nonbasic Humean chances. Then note that (i) Hoefer's basic Humean chances face a serious version of the explanatory challenge. Indeed they face precisely the same version of the explanatory challenge as the actual frequentist must deal with. Since they are nothing over and above the patterns in the events, basic Humean chances cannot explain those patterns. And moreover (ii), we cannot simply assume that the only chances that there are, or that the only chances that there are that play an important explanatory role in our best scientific theories, are nonbasic chances. In particular, the only way in which (2) could be understood as a nonbasic Humean chance would be if Bohmian mechanics turned out to be true. And, as was discussed above, while Bohmian mechanics is one of the live contenders for our best interpretation of quantum mechanical phenomena, it is hardly the only such contender. If, for instance, the GRW interpretation turns out to be correct, then (2) will instead turn out to be a basic Humean chance and, given the explanatory challenge outlined above, will therefore be unable to explain (1).

20. See Lewis (1986, 1994) and Loewer (2004) for a recent extension of Lewis's view. Note that it is possible to separate the best systems analysis from the constraint that all that exists at the fundamental level is the Humean mosaic—see Demarest (2015)—but I will assume that they go together.

In order to extend the BSA to objective probability, one takes all propositions that describe the Humean mosaic and supplements those propositions with an uninterpreted function P(*x*) that assigns to each proposition a number between 0 and 1. One then considers a further theoretical virtue alongside simplicity and strength—the virtue of fit. A system is better fitting to the extent that P(*x*) is higher for true propositions and lower for false propositions, and the best system is the one that best balances simplicity, strength, and fit. The function P(*x*) is then interpreted as the objective probability function.[21] Insofar as we are hoping to give a best systems analysis of nomological probability, the function P(*x*) is interpreted as the nomological probability function.

The BSA clearly avoids the specific version of the explanatory challenge that faced the actual frequency account. On the BSA, nomological probabilities are not actual frequencies, they are magnitudes that play a certain sort of role in the best system. So the advocate of the BSA is not committed to anything explaining itself. But there is still a nearby challenge. This challenge arises because although (on the BSA) probabilities are not equivalent to frequencies, the reason why some probability function ends up being a part of the best system will, at least in paradigm cases, be because of the existence of some corresponding frequency. Insofar as there is a stable, long-run relative frequency in the mosaic, a probability function that assigns a probability close to the relative frequency will in general yield quite a bit of simplicity with little trade-off in terms of fit.

Consider, for instance, the relation between (2) and (1). Although (2) is not equivalent to (1) (according to the BSA), it is still the case that (2) is true because (1) is true—it is because the frequency described in (1) obtains that the probability function that generates (2) is a part of the simplest, strongest, and best-fitting system. And if *p* is true because *q* is true, then surely in some sense *q* explains *p*. But then it follows that (1) explains (2) *and* that (2) explains (1). And this seems to be a worrisome sort of explanatory circularity. How can it be that *q* explains *p and* that *p* explains *q*?[22]

It is important to note that the point generalizes beyond the particular worry about circular explanation. We already established that, according to the Humean, facts about chances must depend on the mosaic, in some sense. More specifically, given the best systems analysis, it seems natural to say that (2) depends on (1). But it is also natural to think that explanatory relations (especially explanatory relations of the type that scientists are after)

21. Here I am following the helpful exposition in Loewer (2004).

22. Is the Humean committed to (1) wholly explaining (2)? Or only to (1) partly explaining (2)? It does not matter for our purposes. The reader who finds the latter more plausible is invited to change each instance of '(1) explains (2)' below to '(1) partly explains (2)'.

track dependence relations. When we introduce (2) to explain (1), then, we are saying that (1) depends on (2). But then, according to the Humean, (2) depends on (1) and (1) depends on (2). And that seems to be a worrisome sort of symmetric dependence. How can it be that *q* depends on *p and* that *p* depends on *q*?

There are two general sorts of moves that an advocate of the best systems analysis might make in response to these challenges. First, they could try to insist that (1) does not explain (2), or that (2) does not depend, in any sense, on (1). But that seems, at the very least, difficult to square with both the Humean motivation and the details of the BSA. Second, they might try to embrace the relevant sorts of explanatory circularity or symmetric dependence and insist that it is just fine for *p* to explain *q* while *q* explains *p* and for *q* to depend on *p* while *p* depends on *q*.[23] It is not, given the argument presented in section 2 of this paper, an option to simply claim that (2) does not explain (1), or that (1) does not, in any sense, depend on (2).

One way to make the second strategy described above (embracing circular explanation and symmetric dependence) more palatable is to follow a suggestion put forward by Loewer (2012) in response to the explanatory challenge that targets advocates of the best systems analysis of laws. According to that challenge, since Humean laws depend on the mosaic—and thus on the instances of those very laws, they cannot in turn explain those instances. Loewer's suggestion is that Humeans (i) distinguish between two types of explanation—metaphysical explanation, on the one hand, and scientific explanation, on the other, and then (ii) claim that the mosaic metaphysically explains the Humean laws, while the laws scientifically explain the mosaic. Extended to Humean analyses of chance, the suggestion would be that (1) explains (2) in the sense that it metaphysically explains (2), while (2) explains (1) in the sense that it scientifically explains (1). And that sort of circular dependence, according to Loewer, is not in fact problematic.[24]

Is this distinction between metaphysical and scientific explanation legitimate? Is it the sort of distinction that a Humean should be happy to accept? Does it suffice as a response to the worry about circular explanation? Can it be extended to ameliorate the worries about symmetric dependence? Perhaps. But even if the answer to all of these questions is yes, there is good reason to be suspicious that the sort of promissory note that the Humeans are issuing here can in fact be cashed. Many of us are operating with a back-

23. For a recent discussion of symmetric dependence relations, see Barnes (forthcoming). Note that while Barnes argues that some dependence relations may hold symmetrically, her argument is entirely compatible with symmetric dependence relations being a significant cost of a theory.

24. Further discussion of this suggestions is found in Lange (2013), Hicks and van Elswyk (2015), and Miller (2015).

ground view on which scientific explanation—at least in paradigm cases—is some kind of causal explanation and on which metaphysical explanation tracks grounding relations, which—whatever else they are—are noncausal.[25] So many of us are operating with a background view on which there is a prima facie reason to think that there is a substantive difference between scientific explanation and metaphysical explanation—one is causal, while the other is not. But it is not at all obvious that this sort of straightforward distinction maps onto the case at hand. In particular, it is not at all obvious that the Humean should think that nomological probabilities cause frequencies in any straightforward sense. Perhaps such a view is plausible on a propensity analysis, according to which probabilities just are causal dispositions or tendencies. But can a Humean tell a similar story? Should she want to?

This point is by no means decisive, but it does serve to shift the burden on to those who wish to defend something like Loewer's move. Perhaps one could say that nomological probabilities mediate causal relations in the mosaic in some more sophisticated sense. Perhaps we can tell some alternative story of scientific explanation that makes sense of the relation between (2) and (1) and maintains a robust distinction between scientific and metaphysical explanation. But until we see that sort of account, we ought to be skeptical of Loewer's way of accepting circles of explanation, and of any analogous attempt to respond to worries about symmetric dependence.

In any case, my goal here is not to argue that there are decisive reasons for thinking that any successful analysis of nomological probability must be non-Humean. My goal is to show that those who wish to give a Humean analysis of nomological probability have their work cut out for them and to emphasize that such work is not work that can be postponed or set aside. The argument in sections 2 and 3 of this article shows that we should think that there are nomological probabilities pretty much regardless of the costs associated with accepting such entities into our metaphysics. Insofar as you are a committed Humean, then, you cannot simply dismiss the explanatory challenge outlined here. You must meet it head on. Until you have done so, your position is untenable.

5. Conclusion. In sections 2 and 3, I argued that insofar as we are naturalists—insofar as we take standard scientific methodology as a good guide to successful inquiry into what the world is like—we ought to believe in the existence of nomological probabilities. This result holds even if it turns out that we cannot give a Humean analysis of such probabilities or if the only available analysis of such probabilities appeals to entities that are supposed to be metaphysically weird in some other sense. Indeed the result holds even if it turns out that we cannot give any robust metaphysical analysis of such

25. Indeed this is precisely how Loewer (2012, 131) presents the distinction.

probabilities at all and instead must accept them as sui generis entities. Scientists are willing to accept metaphysically strange and novel entities into their theories when the alternative is leaving a robust pattern unexplained. Metaphysicians should be too.

In the second part (sec. 4), I presented a challenge for one common group of metaphysical analyses of objective probability—Humean analyses—understood as candidate metaphysical analyses of nomological probability. Perhaps this challenge can be met, but at the very least advocates of Humean analyses of nomological probabilities have quite a bit of work ahead. And insofar as that work cannot be completed we ought to take seriously the possibility that we will need to accept an alternative analysis of nomological probabilities.

The reader will note that I have not attempted to answer the following question: insofar as a Humean analysis of nomological probability fails, what sort of non-Humean analysis should we give? A plausible initial starting point for such an analysis would, of course, be various propensity analyses. But the point I am making here is a more general one. Even if you think that various propensity analyses ultimately fail,[26] you should not give up on the notion of nomological probability. Even if you think that there is no successful analysis of nomological probabilities at all, and that they are instead wholly sui generis entities, you should still think that they exist.

By way of conclusion, let me say something about the particular methodology I have employed above. It can sometimes seem like metaphysicians who are interested in concepts like probability, lawhood, time, or space—concepts that clearly arise in our best scientific theories—face a choice between taking everything that those theories say at face value or being scientifically revisionary. But as the above argument demonstrates, there is a way of doing scientifically respectable metaphysics that goes beyond the former but falls short of the latter. On this way of doing metaphysics we take standard scientific practice as a model and paradigm instance of successful inquiry into what the world is like, while leaving room for the possibility that the particular theories that scientists endorse might be problematic.

In this paper I have argued that in scientific theorizing it is standard practice to be more certain that there are things that play certain types of explanatory roles—in particular that there are things that explain certain sorts of patterns—than of any particular metaphysical constraints on what sorts of things there are in the world. The same should be true of metaphysical theorizing. And insofar as it is, our metaphysics ought to include nomological probabilities.

26. You might think this, for instance, due to the reasons discussed in Humphreys (1985) or Eagle (2004).

REFERENCES

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

Armstrong, David Malet. 1985. *What Is a Law of Nature?* Cambridge: Cambridge University Press.

Barnes, Elizabeth. Forthcoming. "Symmetric Dependence." In *Reality and Its Structure*, ed. Ricki Bliss and Graham Priest. Oxford: Oxford University Press.

Demarest, Heather. 2015. "Powerful Properties, Powerless Laws." In *Putting Powers to Work: Causal Powers in Contemporary Metaphysics*, ed. Jonathan Jacobs. Oxford: Oxford University Press.

Eagle, Antony. 2004. "Twenty-One Arguments against Propensity Analyses of Probability." *Erkenntnis* 60:371–416.

Emery, Nina. 2015. "Chance, Possibility, and Explanation." *British Journal for the Philosophy of Science* 66 (1): 95–120.

Gillies, Donald. 2000. *Philosophical Theories of Probability*. New York: Routledge.

Hájek, Alan. 1996. "'Mises Redux'—Redux: Fifteen Arguments against Finite Frequentism." *Erkenntnis* 45 (2–3): 209–27.

Hicks, Michael Townsen, and Peter van Elswyk. 2015. "Humean Laws and Circular Explanation." *Philosophical Studies* 172 (2): 433–43.

Hoefer, Carl. 2007. "The Third Way on Objective Probability: A Sceptic's Guide to Objective Chance." *Mind* 116 (463): 549–96.

Humphreys, Paul. 1985. "Why Propensities Cannot Be Probabilities." *Philosophical Review* 94 (4): 557–70.

Lange, Marc. 2013. "Grounding, Scientific Explanation, and Humean Laws." *Philosophical Studies* 164 (1): 255–61.

Lewis, David. 1986. *Philosophical Papers*. Vol. 2. Oxford: Oxford University Press.

———. 1994. "Humean Supervenience Debugged." *Mind* 103 (412): 473–90.

Loewer, Barry. 2001. "Determinism and Chance." *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics* 32 (4): 609–20.

———. 2004. "David Lewis's Humean Theory of Objective Chance." *Philosophy of Science* 71 (5): 1115–25.

———. 2012. "Two Accounts of Laws and Time." *Philosophical Studies* 160 (1): 115–37.

Maudlin, Tim. 2007a. *The Metaphysics within Physics*. Oxford: Oxford University Press.

———. 2007b. "What Could Be Objective about Probabilities?" *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics* 38 (2): 275–91.

———. 2013. "The Nature of the Quantum State." In *The Wave Function: Essays on the Metaphysics of Quantum Mechanics*, ed. Alyssa Ney and David Albert, 126–53. Oxford: Oxford University Press.

Meacham, Christopher J. G. 2010a. "Contemporary Approaches to Statistical Mechanical Probabilities: A Critical Commentary—Part I: The Indifference Approach." *Philosophy Compass* 5 (12): 1116–26.

———. 2010b. "Contemporary Approaches to Statistical Mechanical Probabilities: A Critical Commentary—Part II: The Regularity Approach." *Philosophy Compass* 5 (12): 1127–36.

Miller, David. 1995. "Propensities and Indeterminism." *Royal Institute of Philosophy Supplement* 39:121–47.

Miller, Elizabeth. 2015. "Humean Scientific Explanation." *Philosophical Studies* 172 (5): 1311–32.

Popper, Karl R. 1959. "The Propensity Interpretation of Probability." *British Journal for the Philosophy of Science* 10 (37): 25–42.

Sober, Elliott. 2010. "Evolutionary Theory and the Reality of Macro-Probabilities." In *The Place of Probability in Science*, ed. Ellery Eells and J. H. Fetzer, 133–61. New York: Springer.

Suárez, Mauricio. 2013. "Propensities and Pragmatism." *Journal of Philosophy* 110 (2): 61–92.

Williamson, Jon. 2010. *In Defence of Objective Bayesianism*. Oxford: Oxford University Press.