Metaphysical Probability*

Stephan Leuenberger

Many philosophers are realists about modality. They accept that there is a fact of the matter whether a given proposition is possible. While they acknowledge that the term 'possible' is context-sensitive, they insist that it has a privileged, default extension, which in any context is picked out by 'metaphysically possible'. Moreover, they deny that metaphysical possibility is the same as nomological possibility, i.e. compatibility with the actual laws of nature. (They say the same, *mutatis mutandis*, of metaphysical necessity.) Many of these would not make their realism hostage to the availability of an analysis. They may put little stock in attempts to reduce metaphysical possibilitiy to semantic or epistemic facts, such as facts about consistency and ideal conceivability. I call them *non-reductive realists* about modality.[1]

It may be exaggerated to call it "orthodox," but non-reductive realism about modality is currently regarded as respectable. In contrast, non-reductive realism does not seem to be a recognized view in the philosophy of probability. Typically, philosophers reject realism in my sense: They either deny that there is a fact of the matter about whether a proposition has a given degree of probability, or take those facts to be entailed by the laws of nature, or some established scientific theory. In so far as philosophers are realists, they are reductionists, taking probability to be reductively analysable, perhaps in terms of frequencies, credences, or symmetries.

Non-reductive realism about probability appears to occupy an exotic, uncharted region of logical space. In this paper, my aim is to put it on the map, and advertise its attractions. I argue that it should be congenial at least to those who are already non-reductive realists about possibility and necessity.

*Acknowledgments to be added.

1 Probability and Its Interpretations

Probability has been the subject of mathematical inquiry since the middle of the seventeenth century. In the twentieth century, rules for codifying reasoning with probabilities were proposed. The axiomatic system of Kolmogorov proved most influential. It takes a probability function P to be a mapping from a field F of subsets of a given set Ω to the real unit interval. The mapping needs to satisfy the following three axioms: $P(\Omega) = 1$, $P(\emptyset) = 0$, and $P(X \cup Y) = P(X) + P(Y)$ for disjoint X and Y. We may take the members of F to be propositions, construed as sets of possible worlds. The axioms can then be translated as follows: The necessary proposition has probability 1, the impossible proposition has probability 0, and the probability of the disjunction of incompatible propositions is the sum of the probabilities of these disjuncts.[2] From this, the laws of probability calculus are derivable.

There are other formal systems for probability. They may add further axioms to the three above, e.g. countable additivity or regularity. They may be formulated in a different language than Kolmogorov's, say with a two-place function symbol for conditional probability rather than a one-place function symbol, or with a relation symbol standing for "at least as probable as." I will not take a stand here on what the right axiomatization of probability is, since much of my discussion is does not turn on that question. I only use Kolmogorov's particular system for purposes of illustration.

Axiomatic systems can be viewed as implicitly defining the primitives they deploy, provided they are true on just one interpretation of these terms. However, it has long been known that axioms for probability are true under different interpretations of the function symbol 'P'. For example, properties as different as frequencies and rational credences satisfy them. Therefore, such a system cannot provide an implicit definition of 'P'. The formal system does not interpret itself, as it were, and is in need of interpretation by its users.

For that reason, philosophers of probability often talk about there being different "interpretations" of probability.[3] Canonically, five such interpretations are distinguished: the classical, logical, frequency, propensity, and subjective interpretation [Hájek, 2007]. (Sometimes the classical interpretation is omitted, since it is either considered obsolete or a special case of the logical interpretation.) The non-reductive realist about probability cannot be happy with any of them. We may thus take her to advocate a different, "metaphysical" interpretation.

On this proposal, probability is irreducible, it belongs to the metaphysical ground-floor, as it were. It is a graded property of propositions in the same family as the more familiar metaphysical possibility. Probability is graded possibility.[4]

Besides probability, there is another graded or at least comparative concept of possibility, "closeness" or similarity to the actual world, which is used in standard semantics for counterfactual conditionals. To prevent misunderstanding, I would like to point out that "probability is graded possibility" should not be understood as the claim that propositions true in "close" worlds are more probable than those only true in "remote" worlds. It is true that from a closeness comparison for worlds, we could derive a closeness comparison for propositions. (For example, by defining A to be closer to @ than B if some world in which A is true is closer to @ than any world in which B is true.) But that other comparative notion of possibility is independent from probability. The proposition true only at @ is close but highly improbable. A proposition true at all and only the worlds that are remote from @ (on some resolution of the vagueness in 'remote') is remote, but has a relatively high probability. The necessary proposition is close and maximally probable, while the impossible one is remote and maximally improbable. In my usage, 'graded possibility' is always related to probability, not to distance between worlds.

Here is a brief preview of the paper. In section 2, I ask what, if any, concept of probability we deploy in ordinary thinking. I offer some examples that are intended to make it plausible that we do have such a concept, and that it plays an important role in our cognitive life. The examples are designed to bring out the contrast between probability, properly construed, and two other notions that have been proposed as substitutes or analyses: frequency and credence, or degree of belief. In section 3, I briefly rehearse reasons familiar from the literature why the concept we deploy cannot be understood as "logical probability." In section 4, I elaborate on how my proposal ought to be understood. I discuss the time-relative notion of objective chance, and further draw on the analogy to metaphysical possibility. Section 5 asks what reasons we have for embracing realism about probability. If my case in section 2 is successful, we do have the concept of objective probability. But is there something in the world that answers to the concept, and if so, what is it? There is room for an error theory, according to which our ascriptions of probabilities are wrong, not just because they falsely estimate the probabilities, but because there are no such probabilities. While I am not aware of many explicit defenses of such an error theory, I suspect that it is wide-spread among philosophers. I do not have a direct argument against it. Rather, my case for realism, and thus against the error theory, will depend on whether I can make my positive view attractive. I advance a so-called "theoretical utility argument": metaphysical probability needs to be invoked, or is at least usefully invoked, in promising, independently plausible philosophical accounts of the laws of nature, the foundations of intentionality, and of the direction of time.

2 The Concept of Probability

"Interpretation of probability" can mean different things. On one understanding, every concept that is usefully modelled by a certain class of mathematical object, a "probability function" or "(normalized) measure," thereby counts as an interpretation of probability. Since we can represent frequencies and credences by such functions, they both are such interpretations, whether or not they have anything in common beyond sharing a certain mathematically salient structure.[5] On another understanding, an interpretation of probability is a philosophical account of the pre-theoretically grasped concept of probability. Obviously, frequencies and credences are different enough from each other that not both of them can provide an interpretation of probability in that sense.

In this section, I want to present frequentism and subjectivism as foils to non-reductive realism. Rather than telling us what probability is, they change the topic. (Of course, this is not to say that their topics are uninteresting.) Bringing out the sense in which they change the topic will, I hope, shed light on the topic itself. The aim here is not to engage with sophisticated developments of frequency theories and subjectivist theories. While they may be useful for many theoretical purposes, they do not explicate the concept of probability that I am concerned with here.

When philosophers write about probability, their examples tend to involve either the weather or the casino. With these sorts of examples, the distinction between frequencies and subjective on the one hand and genuine probabilities on the other hand can be subtle. I thus introduce different examples, where it is clear that the probabilities involved are single-case and objective, and where the frequentist and the subjectivist interpretation are blatantly inadequate.

The first story serves to emphasize that objective probability is not frequency:

> **The Marriage Proposal** Fred thinks that Suzy is the love of his life. For him, she is unique, there just isn't a woman like her. He intends to make a marriage proposal to her. What, he wonders, is the chance that she will accept? After a great deal of worrying and reflecting, he concludes that the chance is about even.

In this example, Fred is wondering about the probability of a certain proposition, namely, the proposition that Suzy accepts Fred's marriage proposal. This is a single-case objective probability, or objective chance.

Consider the suggestion that Fred is instead wondering about a certain frequency; perhaps the frequency with which marriage proposals get accepted. It is implausible on the face of it that Fred is interested in that frequency. But more importantly, it is part of the story that Fred thinks that Suzy is unique. We can add that Fred thinks that he himself too is unique in his way. Then Fred has to conclude that the frequency in question is either 0 or 1, depending on whether she actually goes on to marry him. But it is part of a story that he think the chance is about even. Thus he would have incoherent attitudes if he took chances to be frequencies. If we add to the story that Fred is a professional logician with coherent attitudes, we get the result that Fred is not wondering about frequencies.

In general, a single-case probability is a property of a proposition, or an event, in the language of probability theorists, while frequencies are properties of event types. Thus frequencies are not even of the right category to be candidates for being chances.

It could be suggested that while Fred is not wondering about a frequency, he is not wondering about an objective probability either. Rather than trying to form a belief about an objective chance, he is merely setting setting his subjective probability for Suzy's marrying him. But that suggestion is wrong too. We can see why by continuing the previous story.

> **Ten Years Later** Fred wonders what the chance was, ten years ago, that Suzy would accept his proposal. Was there in fact a significant chance, or maybe not, because she was already in love with Billy, whom

> she went on to marry? Fred concludes that the chance was more than zero, but less than a half.

In this story, we simply cannot make sense of Fred's attitude if we take the probability involved to be subjective. Fred remembers perfectly well what his subjective probability for Suzy accepting his proposal was ten years ago: it was about one half. Moreover, he is perfectly aware that right now, his subjective probability for Suzy having accepted his proposal ten years ago is zero. Thus he has a belief about a past single-case chance, not just a graded belief about some non-chancy fact.

These examples can be seen as a variation on a theme familiar from the discussion of modality: modal facts are something we care about, they are of concern to us. Kripke [1980] presents the much-cited example of Humphrey, who cares about the possibility that he might have won the election. Such concern is supposed to be hard to account for by theories of modality that are reductive in some sense, such as fictionalism [Rosen, 1990] and, for *de re* modality, counterpart theory [Lewis, 1968]. Similarly, I am using it here against reductive accounts of probability, such as frequentism and subjectivism.

Still, the two examples above leave do not seal my case. You might object now that while Fred is not wondering about frequencies or degrees of belief, he is not wondering about objective probabilities either. The parts of the story where he tries to assign values, "about one-half," or "more than zero, but less than a half," are unrealistic. In our ordinary thinking, you might say, we are not concerned with such numerical values, but only with whether something is, or was, possible. Fred ought to be taken to wonder about possibility, not probability. To deal with this objection, I introduce a third story, designed to show that at least sometimes, our concern with possibilities is proportional to their probability.

> **The Prison Escape** Ken is released from prison after twenty years, regretting that he has missed out on so much in his life. He studies a book on quantum mechanics, and learns that it is possible, according to the theory, that a person walks through a wall. He remembers his first day in prison twenty years ago, when he was constantly banging his head against the wall. He thinks: "It might have happened that I just walked through the wall, and then I would have spent the twenty years in liberty." But then Ken does some calculations, and learns that this chance was minute, so minute that even if there were billions of

> universes like ours, one should not expect it to happen to even one person. After realizing this, Ken's attitude changes. He no longer regrets that it did not happen.

The point about the story is that Ken's attitudes, and how they change when he learns more about the value of the probability, seem very reasonable. Even though his walking through the door nomologically and therefore metaphysically possible, it is clear that given how small its probability was, a feeling of regret about the failure of that possibility to materialize would be out of place.

Let me be clear about what I take these examples to show, and what they clearly do not show. In my view, they bring out that we do have a concept of single-case objective probability, distinct from frequency and credence. We ought to interpret Fred's and Ken's thoughts as deploying that concept. To establish non-reductive realism as characterized, two more things would need to be shown. First, that there is something in the world that answers our concept of probability, that there really are fact of the matter about probability. Second, that the pertinent probabilities are not always those entailed by the laws of nature. (The sense in which the view is non-reductive will become clearer in the next two sections.)

My case for realism will appeal to the theoretical utility of probability. One lesson to draw from the argument to be given in section 5 is that if we were error theorist about probability, we would arguably have to be error theorists about other things as well. As to the second point, the utility argument will invoke probabilities that are not entailed by the laws of nature. Even ideal scientific theories would not be rich enough sources of probabilities to cover our theoretical needs.

3 Reductive Realism?

The frequentist and subjectivist interpretation spell out a useful concept, but not the concept of probability. In contrast, the classical and logical interpretation, in their crude versions, do suggest an analysis of probability, but they fail to spell out a coherent concept. I briefly introduce them here as foils to the non-reductive view. Again, I will be painting with a broad brush. Perhaps sophisticated variants of these interpretations would not be vulnerable to the standard objections rehearsed here, but I doubt that they would keep their reductive realist flavor.

Consider a naive answer to the question what probability is, inspired by Laplace's classical definition of probability, and still widely used in secondary school mathematics:

Probability is the ratio of the favourable to the possible cases.

Or more precisely, the probability of X is the ratio of the cases in which X to all the possible cases. In school mathematics, this is applied to the following sort of problem: What is the probability that the combined result from throwing two dices is 8? The answer is $\frac{5}{36}$, because there are five ways of getting eight from two dice, and there are 36 possible cases.

It is tempting to generalize this, taking "cases" to be worlds. (Here as elsewhere in this article, nothing turns on whether possible worlds are understood as concrete or not.) Then the probability of a proposition is the number of possible worlds in which it is true divided by the number of all possible worlds. A straightforward way to define conditional probability would then be this: the probability of A given B is the number of worlds where the conjunction of A and B is true divided by the number of worlds where B is true. In the above example, B is the proposition that the two dice are thrown, and A the proposition that the combined result from the two dice is 8.

We could try to apply this to Fred. At first shot, the objective probability of Suzy's accepting his marriage proposal is the number of worlds in which she does accept divided by the total number of possible worlds. Or a bit better, the number of worlds in which she accepts divided by the number of worlds in which the proposal is made. Or even better, the number of worlds in which she accepts and which share the past with our world divided by the number of worlds which share the past with our world. (4.1, below, discusses the restriction to worlds sharing the past.)

Unfortunately, this naive proposal does not work, for a familiar reason. If we assume that there is a class of metaphysically possible worlds, then we need to assume that there are infinitely many of them. It would be extremely implausible to claim that there are only finitely many possible worlds. There are infinitely many worlds even among those with just a single particle, since there are there are infinitely many different trajectories that the particle might take. Ratios with an infinite denominator are not in general defined, and hence the naive proposal fails.

Nonetheless, it will become clear in the next section that non-reductive realism preserves something of the spirit of the naive classical theory.

The term "logical probability" may suggest that there is a kind of probability, namely logical probability, to be distinguished from other kinds. But we need not understand the theories in the tradition of Keynes, Jeffreys, and Carnap as positing logical probabilities in that sense. Rather, we could take such theories as trying to give us *a priori* access to probabilities by exploiting logical symmetries, or so-called "principles of indifference." The probabilities whose true values they try to gauge might be those that Fred was wondering about. "Logical" would then not characterize the kind of probability, but the kind of epistemic access we have to it. (I am not here concerned with the intentions of the historical figures associated with the logical interpretation, but with the question whether the techniques they used could vindicate a form of reductive realism.) However, for reasons familiar from the literature on probability, the program of logical probability does not succeed.

Consider a version of what is called "Bertrand's Paradox," a stock example for the sort of problems the program faces, A factory produces cubes whose edges are all between 0 and 2 cm in length. Now it seems that the logical probability that a randomly drawn cube has an edge length of less than 1 cm should be 0.5; the probability that the area of a side is smaller than 1 cm^2 should be 0.25; and the probability that its volume is smaller than 1 cm^3 should be 0.125. However, the propositions that a cube has an edge length smaller than 4 cm, that it has a side area smaller than 1 cm^2, and that it has a volume of less than 1 cm^3 are true in exactly the same possible worlds, and cannot consistently be assigned different probabilities.

Another well-known problem for the attempt to derive probabilities from logical symmetries is the relativity to language. What probabilities are obtained is sensitive to apparently arbitrary choices of a language. However, probability, in the objective sense I am concerned with, is mind-independent, and thus language-independent.

These two problems, relativity to language and incompatibility of different principles of indifference, make it very doubtful whether the program of logical probability can yield a reductive account of probability. Of course, I cannot preempt any proposal for reduction. (Nor will I try to specify precisely what would count as a "reduction" in this context. More on this in the next section.) Nonetheless, I think it is worth exploring a non-reductive option.

4 Probability on the Metaphysical Ground-Floor

Roughly, my thesis is that objective probability is real and irreducible, that it belongs to the metaphysical ground-floor. There is a correct answer to the question that Fred was wondering about, whether or not we can know that answer, and whether or not that answer is determined by non-probabilistic facts. In this section, I want to elaborate on what this thesis amounts to.

Let Ω be the set of all possible worlds, and F a field of subsets of Ω, or propositions.[6] Then there are many functions from F to the real unit interval that satisfy the Kolmogorov axioms. In section 1, I called them "probability functions." In mathematics, they are often called "measures" (or "normalized" measures, a qualification that I will omit). It is useful to be able to refer to those functions with a term that does not suggest that they all represent probabilities.

My claim is that among all these measures on F, one is metaphysically privileged, and gives us the probabilities of propositions. Or rather, it gives us the metaphysical probability of propositions, and plays a crucial role in determining various restricted kinds of probabilities that we may be referring to in certain contexts. I will come back to the last point in 4.1 below.

We sometimes metaphorically speak of the "space of possible worlds," and take the worlds as points in that space. Thus we are accustomed to think of the set of possible world as having more structure than just the one given by the membership-relation and the subset-relation. By claiming that there is a metaphysically privileged measure μ, I am positing additional structure on that set Ω. We can illustrate that structure if we take the metaphor of a "space" literally for a moment. Then $\mu(X)$ is the volume of the subset X of Ω. For we can think heuristically of the measure of a set of points in a space as the volume of the region occupied by the points in the set. To elaborate on this picturesque explanation further, we can think of modal space as having n dimensions, one for every pair of an individual and a fundamental property. Some property axes only have two values 0 and 1 on them; others take scalar values. Once the distances between different values on the coordinate axes are fixed, we can calculate the volume of an n-dimensional rectangle in an obvious way, and can approximate the volume of a well-behaved set of worlds by a series of rectangles.[7]

The presentation of my proposal in this section so far might be misleading with respect to the order of explanation that I envisage. It might suggest that I want to argue first that there is a metaphysically privileged measure μ, and secondly that μ is picked out by our concept of metaphysical probability. While such a "measure-first" view deserves to be explored, I prefer a "probability-first" view. For a given proposition X, the question what the probability of X is has a privileged, default answer: it is simply the metaphysical probability of X. The answers for different propositions together determine a measure μ on the set of possible worlds. Among all the different measures, μ is metaphysically privileged because it represents probabilities. It can then be invoked to serve various theoretical purposes, to be discussed in section 5.

The claim that there is a metaphysically privileged measure needs to be qualified, as I noted briefly in section 1. I have taken it for granted in this section so far that probability has the structure of a normalized measure on the space of possible worlds, that is, a function from Ω to the real unit interval satisfying the Kolmogorov axioms. However, it is no commitment of non-reductive realism that such measures are the correct mathematical representation of probability. There are other systems, besides Kolmogorov's, that have been proposed.[8] This raises a host of interesting questions about the structure of probbility. If it determines an ordering of propositions, what properties does that ordering have? Do probabilities come in numerical grades? If they do, are these grades adequately represented by standard real numbers? It is a familiar problem with standard probability functions that the operation of conditionalizing on a proposition with zero probability is not defined. To anticipate from 4.1, one may wish to obtain the objective chance of a proposition at a world w and time t by conditionalizing on the past history of w at t, and it is plausible that a particular stretch of history of a world has probability zero of occurring. There is thus pressure to assume that metaphysical probability has more structure than is guaranteed by Kolmogorov's axioms.[9] In this paper, I want to steer clear of such questions. One can be realist about some subject-matter without offering a theory that answers all questions about it. For convenience, I will continue to assume that probability does indeed determine a measure.

I have tried to convey a sense of what realism about probability consists in.

What does the non-reductionism of the view amount to? I cannot offer a precise criterion for what would constitute a reduction. But no cut-and-dried criterion is needed, as long as it is clear what the spirit of non-reductionism is. Roughly, an account is reductive if it explains probability in terms of something that we believe in independently of our probabilistic beliefs. If adequate, an account in terms of frequencies or credences would be reductive. On the other hand, an account that derives the metaphysical probability function μ from another mathematical structure, say a metric d the set of possible worlds, might not be reductive in the pertinent sense.[10] For the belief that d is metaphysically significant might be motivated partly by the realization that it allows us to define μ (as well as, perhaps, its utility for a theory of counterfactual conditionals).[11]

The question about the modal status of metaphysical probability has so far been left untouched. Most of the discussion in this paper is not sensitive to whether true claims of the form "the metaphysical probability of X is x" are true in all possible worlds or not. (In contrast, claims about nomological probability and objective chance at a time will clearly be metaphysically contingent.) By default, they should be taken to be true in all worlds, it seems to be, just as claims about metaphysical possibility standardly are.[12] For the assumption that μ_w differs for different worlds w appears to be less theoretically parsimonious. Presumably, metaphysical probabilities at a given world do not just float free of the history of categorical, non-modal facts at a world, but rather supervene on that history.[13] If they do supervene, there is a function f that maps histories to probabilities. That function f itself will not vary from world to world. Thus we would have to postulate world-invariant structure as well as the world-relative structure μ_w, and end up with a less economical theory.

4.1 Chance and Historical Possibility

Construing probability as a sort of graded possibility allows us to transfer various insights from philosophical thinking about possibility to the topic of probability.

We have become used to there being several different notions of possibility. Fred's statement, uttered at 11:10, that it is impossible for him to catch the 11:15 train does not mean that it is metaphysically impossible. It is not falsified by a metaphysically possible world in which a helicopter picks him up and drops him at the train station. As it is often put, the world quantifier is restricted, only a subset of the metaphysical possibilities is under consideration.

The restricted notions of possibility have probabilistic analogues. Fred's statement that the probability of the coin landing head and of its landing tails are not falsified by the non-zero metaphysical probability that the coin lands on the edge, or flies upwards or disappears, or less fancifully, ends up not being thrown. Worlds in which there is no outcome of the coin toss are set aside in that context, just as worlds where Fred has different means of transport at his disposal were set aside in the previous example. The statement about the coin toss implicitly deploys the concept of conditional probability: it assigns a probability conditional on there being an outcome of the toss.

On many uses, 'possible' stands for what we might call "historical possibility," a time-relative notion. Arguably, it was possible in 1750 but is not possible now to prevent the melting of large parts of alpine glaciers. The graded cousin of historical possibility I call 'chance', rather than 'historical probability'. I have so far largely avoided that term, even though for many philosophers, objective probability just is chance. I prefer to reserve 'chance' for the time-relative version of probability.

Chances depend on past history. The chance of rain on Saturday depends on how much water has accumulated in by the end of Friday, and may be different on Thursday from what it is on Wednesday. Chance can be defined in terms of the sort of objective probability that is not relative to time. Indeed, such a definition has been proposed to give a semantics for the chance-operator [van Fraassen, 1980a]. In the following, μ stands for metaphysical probability, the privileged measure on the space of worlds, and H_{wt} the history of world w leading up to time t.

H_{wt}: the history of world w up to time t.

$MC_{wt}(X)$: the (metaphysical) chance of proposition X in world w at time t.

$MC_{wt}(X) = \mu(X/H_{wt})$[14]

Thus the chance of proposition X at time t in world w is $\mu(X/H_{tw})$, the objective probability of X conditional on the history.[15]

: There can only be undermining futures if chance does not supervene on past history. On my account, chance does supervene on past history; hence the argument for the incompatibility does not go through. (It is true, though, that some restricted notions of chance might conflict with the principal principle (if the restriction does not supervene on past history); but that need not be a problem, since the principle is meant to be definitive of absolute chance.)

The history of the world is here meant to be the total history, including the trajectory of every single particle in the universe. What the probability is conditional on is a proposition, which is a set of worlds in the framework adopted here. To make sense of the expression 'the history of the world up to time t,' or 'H_{tw}', it is useful to think of worlds not not as points in a space, but as branches in a tree-structure, as is customary in the semantics of tense logic.[16] H_{tw} is then the set of worlds which share the initial segment with w up to t, and possibly branch out at any time after that.[17]

Objective chance at time t is one example of a kind of probability obtained from conditionalizing metaphysical probability on a certain proposition. Another example is what I call "nomic" or "nomological" chance. It is a standard view that metaphysical possibility is not the same as nomic possibility, the kind of possibility science is concerned with. That something travels with superluminal velocity is a metaphysical, but not a nomological possibility. I submit that the same is true of metaphysical chance and nomic chance, the kind of chance science tries to specify. Nomic chance is naturally construed as metaphysical chance conditionalized on the laws of nature:

$NC_{wt}(X)$: the nomic chance of proposition X in world w at time t.

N_w: the laws of nature of w.

$NC_{wt}(X) = \mu(X/H_{wt} \wedge N_w)$

The plausibility of the claim that metaphysical and nomic chances differ depends on one's view of law of nature. If what the laws are does not supervene on an initial segment of the history of the world, then for various mutually incompatible candidate law systems, it is still metaphysically possible at a given time that they are the laws. Then the metaphysical chance at t is not conditional on the laws. According to the regularity theory and its more sophisticated descendant, the Best-System Analysis (to be discussed in 5.1), what the laws are only supervenes on total history, not on some initial segment up to time t. It is also consistent with, though perhaps not entailed by, certain other metaphysical theories of laws that it is not settled at a given time what the laws are.

One lesson from the philosophy of modality is that a plethora of distinct notions may be systematically related in a unifying framework, according to which they are restrictions of one absolute notion. Analogously, we can systematically relate various notions of probability as conditionalizations of one absolute notion. Moreover, when we talk about possibility, it is often not just difficult to know whether we speak truly or not. Often, it appears indeterminate whether we do, that there is no fact of the matter, and hence that there is nothing to know. If Ken asserts that it is possible to fly from London to Sydney non-stop, Fred might be unsure how to evaluate this claim. Presumably, it has not been done, and normal mortals cannot do it at present since there are no commercial flights available. Surely, however, the technology to do it is available.

There is a standard story about this sort of indeterminacy. It is not that the modal facts are indeterminate, it is just indeterminate what restriction was in place when the sentence was uttered. The context may narrow down what restrictions there may be, but it often does not determine it uniquely.

For probability, the indeterminacy problem seems to loom even larger. It is implausible that when we ask about the chance of rain tomorrow, there is a fact of the matter about which particular real number would provide the answer to our question. But the diagnosis of the apparent indeterminacy carries over from the discussion of possibility. The context of the question may not fix what proposition, exactly, μ is supposed to be conditionalized on. Is it just the past? Is it the past plus certain fundamental laws? Or the past plus high-level meteorological laws? Or the past and both fundamental and high-level laws? If the question does not have a determinate answer, it is because it is not determinate which conditionalization is demanded in the context in which it was uttered.

4.2 The Analogy between Possibility and Probability

Attractive and widely accepted theses about possibility correspond to unjustly neglected theses about probability. To sum um the analogy that I have drawn, I first present a list of claims about possibility, and then a list of corresponding claims the non-reductive realist about probability advocates.

1. For every world and every proposition, whether the proposition is (metaphysically or unrestrictedly) possible at the world is an absolute, determinate fact.

2. Possibility cannot be reduced to semantic or epistemic facts.

3. Besides unrestricted possibility, there are various restricted notions of possibility.

4. Many sentences with a possibility-operator are indeterminate in truth-value because the context does not determine the restriction sufficiently narrowly.

5. The possibility talked about in scientific language is some restricted possibility.

For philosophers subscribing to the above views, the corresponding views about objective probability should be worthy of consideration:

1.* For every world and every proposition, what (metaphysical or unconditional) probability the proposition has at the world is an absolute, determinate fact.

2.* Probability cannot be reduced to semantic or epistemic facts.

3.* Besides metaphysical probability, there are various conditionalized notions of probability.

4.* Many sentences about objective probability are indeterminate in truth-value because the context does not determine the conditionalization sufficiently narrowly.

5.* The probability talked about in scientific language is some conditionalized probability.

5 A Theoretical Utility Argument

My argument for accepting probability on the metaphysical ground-floor appeals to the theoretical utility of doing so. The paradigm theoretical utility argument, due to Quine and Putnam, tried to establish the existence of sets, or mathematical entities more generally. (Often, this goes under the label "indispensability argument.") In metaphysics, there are at least two well-known recent examples of arguments of that type. First, there is the case in Lewis [1986a] for the existence of possible worlds. Secondly, there are arguments in Armstrong [1978] and Lewis

[1983a]) for a metaphysically significant distinction between sparse and abundant properties.

Such utility argument try to show that one assumption serves various different and important theoretical purposes. Sometimes, this is put by saying that an entity can do various jobs. The last section already indicated how time-relative objective chance can be construed as a function of metaphysical probability. The so-called "argument from fine-tuning" is an example of an argument that appears to presuppose a probability distribution which is not relative to time. It is claimed that if the values of the fundamental constants were only slightly different, there would be no life on earth. According to proponents of the argument, the truth of the extremely improbable, but in some sense good proposition that there is life provides evidence for a benevolent designer. To spell out that argument, one would need to appeal to a background probability distribution P for the values of the fundamental constants. Moreover, it is hard to see how the argument could work if P is anything less than full-blown metaphysical probability.

However, the argument from fine-tuning has many controversial aspects, and I do not wish to endorse it. Instead, I want to sketch the theoretical work that a metaphysically privileged measure can do in three *prima facie* attractive philosophical accounts: the Best-System Analysis of laws of nature, reference-magnetic descriptivism in foundational semantics, and an account of the direction of time in terms of the asymmetry of convergence and divergence of worlds.

In each case, it will be clear that the probability measure invoked is not specified by the laws of nature. In this respect, the non-reductive metaphysical interpretation sets itself apart from the propensity interpretation of probability, the only member of the canonical five with which it has not yet been contrasted. It is a bit misleading to speak of "the" propensity interpretation, since they form a broad family. The feature that matters for present purposes is shared by all of them, though: probabilities, that is propensities, are the sort of thing that the laws of nature, or ideal scientific theories, need to specify. It is compatible with the non-reductive realist view developed here that some propensity theory gives an adequate account of nomic chance. It cannot offer a story about metaphysical probability, however, and is thus at best incomplete.

5.1 Laws of Nature

A select few of all the truths are laws of nature. What distinction sets them apart from the rest? A promising and influential account is given by the so-called "Best-System Analysis," or "BSA," which is mainly associated with Mill, Ramsey, and Lewis. The rough idea is that laws of nature give a concise, systematic account of what happens in total history. A law of nature is a theorem of the best system of the world.[18]

This raises the question what a system of the world is, and what makes a system better than another. For Lewis, systems are theories formulated in what we may call the "canonical language." He requires that the primitive vocabulary of the canonical language only refers to perfectly natural properties and relations.[19] A system is thus is just a set of sentences, perhaps deductively closed, of that language. The conjunction of the sentences in this set, which may be infinite, expresses a proposition.

What makes a system better than another, with respect to the actual world? There is one obvious sufficient condition: a true system, i.e. one that expresses a proposition that is true in the actual world, is better than a false system. However, there will be many true systems, and we need a standard for comparison among those as well. Lewis proposes that a true system is better than another if it strikes a better balance between simplicity and strength, or informativeness. That account is hard to evaluate until we are told more about how to compare and balance simplicity and strength.

Intuitively, p is simpler than $(p \wedge q) \vee (p \wedge \neg q)$. How we should extrapolate from such cases and assign a simplicity score to systems is not clear. But at least, there are several *prima facie* reasonable options for assigning such a score: we could count the number of symbols, or the number of non-logical primitives, or the number of quantifier alternations, or count all of them and take the simplicity score to be a weighted sum of these. While I cannot discuss whether any of these is ultimately satisfactory, I take it that the notion of simplicity is understood sufficiently well to be used in the BSA.

If system T implies system T', and T' does not imply T, then T is stronger, or more informative, than T'. In that special case, a strength comparison is unproblematic. But there are two reasons to think that the BSA needs more than just comparisons of that type.

The first reason is that strength is to be balanced with simplicity. If T and T' are both true, but T is stronger, then whether T is better overall depends on how much stronger it is than T' (and of course also on how much simpler it is). The BSA apparently does not need just a partial ordering of systems by strength, but something like a strength score which can weighed up with a simplicity score.

The second reason is that competing true systems may be logically independent. Here is a fictional example. In world w, let us grant, some set of candidate dynamical laws contributes so much to simplicity and strength that it is included in any promising system. These dynamical laws are mildly indeterministic towards the future: worlds in which they are true may diverge after sharing an initial segment, but they are not going to differ dramatically. In contrast, these laws are strongly indeterministic towards the past: there is convergence, or near-convergence, of worlds with radically different initial conditions to w. Consider now two candidate systems: T consists of the dynamical laws and a concise, informative generalization about the initial conditions. T' consists of the dynamical laws and a concise, informative generalization of the conditions at some time late in the history of w. Intuitively, T is better than T' since it is more informative. However, the two systems are logically independent, due to the indeterminism of the dynamical laws. If it cannot appeal to a more fine-grained notion of strength than the one considered above, the BSA cannot account for T's superiority over T' in balancing simplicity and strength. (The case can be chosen in such a way that the strength of T more than compensates for its deficit in simplicity in comparison with the system that consists of the dynamical laws only; and that the conjunction of T and T', while stronger than either, pays too high a price in simplicity for its gain in strength.)

Fortunately, the BSA can draw on the idea that the information content of a proposition and its probability are related [Bar-Hillel, 1964, "Semantic Information and its Measures," pp. 298-312]. Provided that we have a measure μ on the space of possible worlds, we can take the strength of a proposition X is simply the measure of its negation, $\neg X$: $strength(X) = \mu(\neg X)$.[20]

A strong proposition is one that rules out many possibilities. Necessary truths, to take the extreme case, are very weak because they do not rule out any possibilities at all. On this proposal, they would have strength zero.

As noted before, there are many probability function, or measures, on the space of possible worlds. This threatens to make it relative to a choice of a measure what the laws of nature are. Intuitively, though, laws obtain mind-independently, and their status ought not to depend on an arbitrary such choice.[21] If there is a metaphysically privileged measure μ, given by the metaphysical probabilities, then the Best-System Analysis can foil the threat of unwanted relativity.

It is one thing to claim that a privileged measure μ allows us to make sense of the notion of strength. It is another thing to claim that it is needed to make sense of that notion, that there is no other solution to the problem of how to balance strength with simplicity. I cannot offer a general argument that this metaphysical assumption is indispensable as well as useful. But some other proposals that come to mind do not work. Certainly, we cannot just count the propositions implied by a system, since there will be infinitely many. The cardinality of the set of implied propositions will be the same for all reasonable candidate systems. It will not do either to simply count how many fundamental facts are implied, where a fundamental fact is the instantiation of a universal, say. Since typical systems will consist of universally quantified sentences, no particular facts will be entailed by them. It thus seems to me that there is at least a good *prima facie* case for the need for a privileged measure in defining strength.

In 4.1, I argued that given a privileged measure μ on the space of worlds, we can take the objective chance of a proposition at a world w and time t to be the result of conditionalizing μ on the history of w up to t. Lewis provides a rival account of objective chance by extending the best-systems. The canonical language in which systems are formulated is now also supposed to contain an uninterpreted term 'chance', and the best system is the one that provides the best combination of simplicity, strength, and fit. The fit of a system at a world w, according to Lewis, is the number x such that the system implies '$chance(H_w) = x$', where H_w is the history of the whole world. This extension of the BSA is problematic, not least because the notion of fit has not been satisfactorily explained.[22] While the BSA is an attractive account of laws of nature, in my view, its extension to chance looks artificial. But for present purposes, the crucial point is this: the rival account of nomic chance likewise relies on the privileged measure on the space of worlds. Thus even if I am wrong in 4.1 about how objective chance relates to μ, I can still run a theoretical utility argument.

5.2 The Foundations of Intentionality

A metaphysically privileged probability measure, I have argued, is needed to cash out informativeness, which is invoked in an attractive account of laws of nature. But informativeness may also play a role in foundations of thought and language. The question how physical entities, such as sequences of sounds, strings of inscriptions, or structures in the brain, can have meaning or intentional content is an old philosophical chestnut. Here, I explain how informativeness, and thus probability, is helpful in spelling out an answer to that question that has been influential lately: global descriptivism.[23]

The term 'descriptivism' is most commonly used for a descriptive semantic theory according to which proper names and certain nouns are synonymous with definite descriptions. Here I use it for a foundational semantic theory which has been brought to prominence by David Lewis. A story much like the one told in Lewis [1970] can be used to motivate it. Suppose a new word is introduced in a linguistic community, for example 'electron'. The term plays a role in some of our theories. It occurs in sentences such as "electrons are parts of atoms," in which other terms (logical vocabulary, 'parts' and 'atoms') are already interpreted. The new term 'electron' is then implicitly defined by those theories. Since there is no sharp line to be drawn between theoretical terms like 'electron' and non-theoretical ones, the story may apply to terms like 'leopard, 'cabbage', and 'cupboard' as well. Descriptivism also allows that several new terms are introduced at once, typically with fairly large theories.

However, in an obvious way the story is inadequate to explain how reference or intentionality come into the world. It only tells us how to use some meaningful items, i.e. the part of the language that is already interpreted, to give meaning to other expressions. Lewis was well aware of this, and suggested that there need not be any previously interpreted words [Lewis, 1984]. Imagine we have a large sentence T, an expression type having a certain structure (types of inscriptions, sound sequences, etc.). Assume that it is determinate how this geometrical or sound pattern should be parsed, i.e. that it has a determinate syntactic structure, and that somehow the logical vocabulary has got its interpretation. Then we consider interpretations, in the model-theoretic sense, of the non-logical terms in T. An interpretation in the model-theoretic sense is here a function assigning to each possible world a domain of individuals, to each individual constant an individual, and to each predicate an intension; the intension assigned to a 2-place predicate, for instance, is a function from worlds to sets of pairs of individuals in that world. We can think of each such intension as representing a property.

The question is now which interpretation correctly represents what the terms in T refer to? A character we may call a "pure global descriptivist" answers as follows: the correct interpretation of T is the one under which T comes out true. Unfortunately, pure global descriptivism is inadequate as a foundational semantic theory. The so-called 'model-theoretic arguments' (Putnam [1980] and Putnam [1981]) can be used to discredit it.

Pure global descriptivism implies that almost no consistent theory is false—the only way for a consistent theory to be false is for it to misjudge how many individuals there are. More precisely, it is only false if none of its interpretations has a domain the size of the set of actual individuals. This is clearly an absurd result: theories can be false in ever so many ways, not just by getting the total the number of actual individuals wrong. The basic idea behind these model-theoretic arguments are familiar, and the technical details do not matter for present purposes.

A response strategy to Putnam's *reductio* of global descriptivism is to narrow down the field of candidates for being the correct interpretation. Taking up a suggestion by Merrill [1980], Lewis (in Lewis [1983a] and Lewis [1984]) imposes a constraint on what classes of possible individuals can serve as intensions of predicates, or are *eligible*, as he puts it. He famously appeals to distinction between natural and unnatural classes. Roughly, a class is natural if its members resemble each other, for example the class of red things. Correspondingly, a class is unnatural if its members do not resemble each other. Typically, the shortest way of picking out such a class is simply by listing its members: Carnap's briefcase, the Eiffel Tower, the number seven, and the planet Jupiter, say.

In a metaphor suggested by Lewis, the word is attracted by natural properties when it is looking for an interpretation in the world. Natural properties serve as "reference magnets." Global descriptivism with an eligibility constraint is not vulnerable to Putnam's *reductio*: A theory might not be true under any eligible interpretation, and thus it might be false. But Lewis went beyond just blocking the model-theoretic argument, and developed further what I call "reference-magnetic descriptivism," a foundational semantic theory to put in place of the discredited pure global descriptivism.

First, he observed that naturalness, and thus eligibility, are a matter of degree. On his picture, everything in any possible world supervenes on the distribution of the perfectly natural properties and relations. At least as it is relevant for reference, naturalness is a matter of definitional distance from perfectly natural properties. He suggests that properties have Boolean structure, and that to be unnatural is to be disjunctive. The grue things form a less natural class than the green things, for example.

In a second step, Lewis replaces truth as a constraint with truthlikeness, which is also matter of degree. He then claims that naturalness (which is not world-relative) and truthlikeness (at some world w) can be balanced. On his account, the correct interpretation of T at w is the one that strikes the best balance between naturalness and truthlikeness at w.

Is reference-magnetic descriptivism a good foundational semantic theory? Such a theory can be tested against our knowledge about the semantic and the non-semantic facts of our world. We know the semantic fact that many of our terms pick out properties of middle-sized objects. We know the non-semantic fact that properties of middle-sized objects are not perfectly natural, but arise from more natural properties of smaller things. If reference-magnetic descriptivism can be shown to imply that none of our terms refer to properties of middle-sized objects, then it is false.

Robert Williams [2007] presents an argument to the effect that Lewis's theory does indeed imply that, or at least that it implies it for some worlds that are macroscopically indistinguishable from ours. Here is a sketch: Model theory tells us that any consistent T is true under some interpretation whose domain consist only of mathematical objects. Accordingly, such an interpretation assigns classes of mathematical objects to the predicates of T. In general, these classes are not natural properties. But Williams points out that under certain conditions, an upper bound can be put on how unnatural they are. He assumes that since T is true under the mathematical interpretation I_M, I_M has a maximal truthlikeness score. He then compares the naturalness score of I_M with the naturalness score of what we know to be the correct interpretation of our theory T in our world—i.e. the interpretation according to which 'table' refers to tables, 'stone' refers to stones, 'green' to green, and so on. To the embarassment of reference-magnetic descriptivism, which one of these two interpretations is more natural turns out to depend on contingent and epistemically open facts about the microstructure of our world.

As mentioned above, Lewis specifies naturalness as definitional distance from the perfectly natural properties. For all we know, properties of quarks may be perfectly natural. But for all we know, the perfectly natural properties are had by entities on an even smaller scale. The more layers of properties of smaller and smaller things there are, the greater the definitional distance from the perfectly natural properties to the properties assigned by the correct interpretation—being a table, being a stone, being green, and the like. Since there is no limit to how many layers there could be, what is in fact the correct interpretation may well score worse on the combination of naturalness and truthlikeness than the mathematical one.

Moreover, even if our world happens not to be among them, there will be possible worlds macroscopically indistinguishable from it that do have that many layers. But it is absurd to claim that 'table' refers to mathematical objects rather than to artifacts in some world that is macroscopically indistinguishable from ours.

The key to a response on behalf of reference-magnetic global descriptivism, I suggest, is to challenge the assumption that any two true theories are equally truthlike. The theory under the mathematical interpretation is true, but it may not be very truthlike. Even if it were more natural than the intended interpretation, it might lose the competition due to its poor truthlikeness score.

In the literature explicitly concerned with truthlikeness, starting from Popper [1963], closeness to truth is typically distinguished from truthlikeness. [Hilpinen, 1976] is particularly clear on this: "The truthlikeness of a proposition depends on two factors: on its closeness to truth, and on the degree to which it conveys information about the truth" [24] If that is right, then T's truth under an interpretation I does not guarantee that it is maximally truthlike under I. Surely, if an interpreted theory is true, it is maximally close to the truth. But the degree of closeness may differ from the degree of truthlikeness.

As Hilpinen notes, truthlikeness is also a function of informativeness. There is a good sense in which it is more truthlike that China has more than one billion people than this it has more than three people, and in which relativity theory (provided it is true) is more truthlike than tautology.

The closeness to truth and the informativeness of a proposition are largely independent. They correspond to the two different graded notions of possibility distinguished in section 1. Whatever is true is maximally close to the truth. Tautologies are true but maximally uninformative. A complete description of the actual world is both maximally close to the truth and maximally informative. A complete description of a world very unlike ours is very far from the truth and maximally informative. Tautologies are maximally close to the truth but minimally informative. The proposition consisting of all and only remote worlds is far from the truth and uninformative. In the practice of asserting, there is often a trade-off between being informative and risking to say something far from the truth.

If truthlikeness is not just closeness to truth, then Williams' argument is blocked as it stands, since there is a factor it does not take into account in the comparison between the mathematical and the correct interpretation. But is there any reason to think that the mathematical interpretation will do badly on the informativeness score? Clearly, there is, since mathematical truths arguably tell us nothing about the world, similar to logical truths. Thus it appears that T under I_M is not very truthlike, albeit maximally close to the truth.

I anticipate the objection that this response is *ad hoc*. Even if there is some notion of truthlikeness that has informativeness as a component, it might be argued, there is no independent reason why that notion ought to be deployed by reference-magnetic descriptivism. However, a case can be made that requiring informativeness is very much in the spirit of that view. One of the motivating ideas of descriptivism, roughly put, is that some connection to truth is constitutive of meaning and representation. Terms get reference in such a way as to make the theories in which they occur true. But if there is a constitutive connection between truth and representation, there ought to be also a constitutive connection between information and representation. There is something *prima facie* odd about a language that can only express logical truths, and thus cannot convey information about the world.

If this is right, a promising version of reference-magnetic descriptivism requires that the correct interpretation for T is the one under which T achieves a good combination of naturalness, closeness to truth, and strength. There may well be further constraints on the comparison of the different interpretations. But what is crucial for present purposes is that the informativeness of a proposition X needs to be invoked, and for the reasons given in 5.1, this should be given by $\mu(\neg X)$, where μ is the metaphysically privileged measure on the set of possible worlds.[25]

5.3 The Direction of Time

The problem of the direction of time supplies another example of the theoretical utility of metaphysical probability. While spatial dimensions have no direction, time does: it flows from the past to the future, as it were. One might simply accept it as an irreducible fact about the world that time displays such an asymmetry, but many philosophers take this to be an unsatisfactory option. The asymmetry of time ought to be grounded in some other asymmetry. A proposal to that effect was made by David Lewis, but it has been questioned whether it succeeds. As I will explain, a privileged probability measure on the space of worlds helps that proposal answer a challenge.

Lewis [1979] suggests that the asymmetry of time is due to an asymmetry of counterfactual dependence, which in turn is due to an asymmetry of convergence and divergence of worlds: it takes much more for worlds to converge than to diverge. If the histories of two possible worlds w and w' match up until time t, but then one particle swerves away from the trajectory prescribed to it by the laws of w, then in general w and w are going to be different at all times after t. In other words, a tiny swerve can make w' diverge from w. But if the histories of w and w' are different before t, then no small violation of the laws of w will ensure that w and w' match after t. It they were to match, all the traces from events in the past of w' that did not happen in w would need to be erased, and traces of events that happened only in w created. In other words, only a large-scale conspiracy can make w' converge to w.

It is widely considered to be a problem for Lewis's proposal that the fundamental physical laws of our world do not seem to underwrite the asymmetry of convergence and divergence, since they are time-reversal-symmetric: if you film some lawful process and then run the film backwards, what is shown is still a lawful process [Elga, 2000]. Indeed, the time-reversal symmetry of laws is not just a problem for Lewis's particular proposal. It is what makes accounting for temporal asymmetry difficult in general.

Fortunately, Lewis's proposal can get off the ground if we can appeal to a metaphysically privileged probability measure. Consider a line that represents the time-coordinate of given world. Any point on the time-axis partitions the history of a world w into two segments; I call them C_t and D_t. Without knowing the direction of time, we can tell whether for different times, two segments are "on the same side," as it were. (To be precise: segments x and y are on the same side if they overlap and their complements overlap too.) By stipulation, all the members of the C-family are on the same side, and all the members of the D-family are on the same side. The task is then to specify what it is for the C-segments to be in the past of the D-segments. I propose the following:

The C-segments are to the past of the D-segments if for all times t, $\mu(C_t) > \mu(D_t)$.

This is meant to capture the intuitive idea that there "more" worlds diverging to w than are converging to it.

This account has the welcome consequence that it is contingent whether a world is temporally asymmetric. Maybe some worlds do not show that sort of asymmetry. The proposal leaves room for modifications; the condition could be weakened by requiring the inequality to hold only for most t, or strengthened by requiring $\mu(C_t)$ to be significantly greater, rather than just greater.

Given this account of the direction of time, we have an example of theoretical unification. Recall the proposal regarding nomic chance $NC_{wt}(A)$ given in 4.2:

$$NC_{wt}(A) = \mu(A/H_{wt} \wedge N_w)$$

(H_{wt} is the history of w up to t, and N_w is the conjunction of the laws of nature of w.) The metaphysically privileged measure μ is used three times here, once explicitly, and twice implicitly, since it goes into both the characterization of the laws of w (assuming that the BSA is largely correct) and of what counts as the past of t. Its triple role in the account of nomic chance testifies to its theoretical utility.

6 Conclusion

There are two main themes in the case for non-reductive realism about probability presented here. First, the view is less outlandish than it might seem, since it shares many features with non-reductive realism about possibility. Secondly, the view promises theoretical benefits.

It needs to be acknowledged that arguments from theoretical utility are hardly irresistible. There are at least three response strategies. First, to dispute that the theories discussed are adequate even if they help themselves to a privileged probability measure. Second, to argue that the theories can overcome their problems without invoking that measure. If these responses fail, there is always the third option of rejecting the theories simply because of their commitment to metaphysical extravagence. I will not attempt to preempt any such moves here. What I hope to have achieved is to put non-reductive realism about probability on the map of theoretical options that deserve to be taken seriously.

I conclude with a bit of propaganda. The metaphysics of modality and probabilistic epistemology (a program sometimes called "Bayesianism") can be counted among the great success stories of philosophy of the last few decades. My proposal for understanding probability is inspired by and draws on both of these. It borrows the basic strategy from probabilistic epistemology, and techniques for its implementation from the metaphysics of modality.

Belief is a central notion in epistemology. Traditionally, beliefs are construed as all-or-nothing: either you believe that p, or you do not believe that p. Much of epistemology is about all-or-nothing beliefs. But recent decades have seen a flourishing of probabilistic epistemology, which takes beliefs to be graded on a scale from 0 to 1. That approach has been able to address many problems more fruitfully than non-probabilistic epistemology. Just as belief is central in epistemology, possibility is central in metaphysics. For almost any X, there is a prominent philosophical account of X that deploys the notion of possibility (or the interdefinable notion of a possible world). One of my main aims in this article has been to make it plausible that just as a probabilistically graded notion of belief is theoretically useful in epistemology, a probabilistically graded notion of possibility can do work in metaphysic.

Even if we replace the ungraded with the graded notion of possibility, we can still draw on the success story of the philosophy of modality. In the past hundred years, different notions that were previously often conflated have been distinguished. As a consequence, many modal arguments have been clarified, and the inferential relations among modal statements have been captured in the powerful and elegant framework of possible world semantics. If we think about probability in the way suggested here, we can transfer many of these lessons learnt.

Notes

[1]A non-reductive realist, in my sense, may not rule out that an attempt at reductive analysis succeeds in the future. She merely denies that realism needs vindication by such an analysis.

[2]In mathematical works on probability, subsets of Ω are typically called "events" rather than "propositions," and elements of Ω "elementary events" rather than "possible worlds."

Some authors define probability functions on a set of sentences closed under Boolean operations, rather than on a field of subsets of a given Ω. I will not consider that approach here.

[3]Alan Hájek [2007] notes that it is misleading to talk of "interpretations of probability," and suggests reading 'interpreting probability' as "providing analyses of various concepts of probability."

[4]I borrow this slogan, though not the theory of probability it is meant to capture, from van Fraassen [1980b].

[5]I am granting for the sake of the argument that frequencies and credences are probabilities in the mathematical sense. While frequencies satisfy the axioms mentioned above, they do not satisfy countable additivity, which is often considered to be part of probability theory. In reality, most people's credences do not satify the probability axioms. Whether rationality requires agents to have probabilistic credences is controversial (see for instance Hájek [Forthcomng]).

[6]I here ignore the technical complications that arise if there are proper-class many possible worlds, as argued in Nolan [1996].

Some philosophers who work with possible worlds do not accept that there is, independently of any conversational context, a definite totality of all possible worlds. (This is a recurring theme in Stalnaker [2003].) For such philosophers, the elaboration of my thesis that I am about to offer will admittedly not be helpful. In this part of the paper, I am not attempting to persuade them, since they are not realists about metaphysical possibility.

[7]The picture can be made more complicated by including relations, and determinables that take structured values such as vectors.

The volume defined is not invariant under non-uniform stretching of coordinate axes (though it is invariant under uniform stretching, i.e. stretching of all axes by the same factor). Thus the distance ratios between pairs of points on different axes are not arbitrary.

I am assuming here that the volume of Ω, the set of all worlds, is bounded.

[8]The relationship between many such systems is usefully investigated in Fine [1973].

[9]It might be represented by a non-standard probability function, or by a conditional probability function, or by a series of lexicographically ordered functions. (For conditions under which these approaches are or are not equivalent, see Halpern [2001].)

[10]In the analogous case of possibility, an account in terms of semantic rules, idealized belief, or polycosmological fictions would count as reductive, but Lewis's realism about possible worlds would not. For Lewis's belief that there are non-actual concrete universes is partly motivated by the realization that they can play a role in a theory of possibility.

[11]A proposal to define a probability measure in terms of a metric was made in Bigelow [1976].

[12]Though Chandler [1976] and Armstrong [1989] deny, for different reasons, that what is possible in this world is possible in all worlds.

[13]I am here working with an intuitive notion of a categorical or non-modal fact and do not presuppose a particular theory about it.

[14]A full account of chance would have to deal with questions that are outside the scope of this article. For example, it would need specify how a time t in w is reidentified in worlds that do not share the initial segment of w up to t, in order to deal with iterations of the chance-operator.

[15]This account of chance is compatible Lewis's Principal Principle, which links chance to rational credence. The so-called "problem of undermining" [Lewis, 1994] does not arise, since chances supervene on past history, not just on total history.

Some authors take it as a condition of adequacy of an account of chance that it entails the Principle Principle, not just that it is compatible with it. It is unclear whether any account has succeeded in doing this (for an interesting proposal, see Hoefer [2007]). In any case, I reject that demand. A theory of chance need not also be a theory of rational credence, and if it is not, it will not entail the Principal Principle. (I also agree with Ned Hall [2004] that the Principal Principle does not capture everything we know about chance.)

[16]In this paper, I ignore complications due to relativity theory for a formulation of how chances depend on the past. To deal with the, one could draw on the technical work of Nuel Belnap and others on branching space-time, e.g. Weiner and Belnap [2006].

[17]This talk of worlds sharing an initial segment and then branching does not have to be taken literally, even under the pretense that the worlds are real and concrete. What is required is only that their initial segment are intrinsic duplicates of each other, without being identical. Branching is then just becoming intrinsically different, what Lewis calls "divergence" [Lewis, 1983a].

[18]Some of the details of how the BSA is best formulated need not detain us here. For example, there is a question whether we should take only the axioms rather than all the theorems as laws, and whether we should take only regularities (presumably, sentences starting with a universal quantifier) as laws. Lewis develops the BSA in Lewis [1973], Lewis [1983a], and Lewis [1994].

[19]There are many interesting question about the canonical language that I cannot take up here: Which mathematical properties and relations count as perfectly natural and are thus referred to by primitive vocabulary? What logical resources are there for defining new terms and predicates from the primitive ones?

[20]That strength is related to exclusion of possibilities has been recognized. For example, Loewer [2004, p. 1119] writes: "Strength is measured in terms of informativeness (possibilities characterized in terms of natural predicates excluded)."

[21]Lewis recognizes this problem, and suggest in response that nature may be "kind to us," such that the best system may be "*robustly* best—so far ahead of its rivals that it will come out first under any standards of simplicity and strength and balance" Lewis [1994, section 3]. Even granting that this pious hope is borne out, it is not clear how it helps the BSA. Surely, an analysis must get the intension and not just the extension of the concept of a law of nature right, and there are possible worlds where no system is robustly best in Lewis's sense. Perhaps Lewis takes the BSA to imply that it is indeterminate what the laws are in such worlds. But then an

argument would be required to convince us that the indeterminacy in the analysans is indeed matched by indeterminacy in the analysandum.

[22]For some of the problems it faces, see Elga [2004].

[23]It is invoked, for example, in Theodore Sider's argument against the view that ontological disputes are merely verbal [Sider, 2001, Introduction].

[24]Lewis [1986a, pp. 24-26] discusses Hilpinen's influential paper.

Hilpinen himself claims that "the number of possible worlds (or a measure of the set of possible worlds) excluded by P is not a good measure of the information conveyed by P" (p. 30), but I find his reasons for doing so are unconvincing. For a criticism of Hilpinen's alternative proposal, see Oddie [2007, section 3].

[25]A full-dress defence of reference-magnetic descriptivism would also need to address the putative counterexamples offered in Hawthorne [2007].

References

David M. Armstrong. *A Combinatorial Theory of Possibility.* Cambridge University Press, 1989.

David M. Armstrong. *Universals and Scientific Realism.* Cambridge University Press, 1978. Volume I and II.

Yehoshua Bar-Hillel. Addison-Wesley, Reading, Mass., 1964.

John C. Bigelow. Possible worlds foundations for probability. *Journal of Philosophical Logic*, 5:299–320, 1976.

Hugh Chandler. Plantinga and the contingently possible. *Analysis*, 36:106–109, 1976.

Adam Elga. Statistical mechanics and the asymmetry of counterfactual dependence. *Philosophy of Science*, supplementary vol. 68:313–324, 2000.

Adam Elga. Infinitesimal chances and the laws of nature. *Australasian Journal of Philosophy*, 82:67–76, 2004. Reprinted in Jackson and Priest 2004.

Terrence L. Fine. *Theories of Probability. An Examination of Foundations.* Academic Press, 1973.

Alan Hájek. Arguments for—or against—probabilism? In Franz Huber and Christoph Schmidt-Petri, editors, *Degrees of Belief.* Forthcomng.

Alan Hájek. Interpretations of probability. *Stanford Encyclopedia of Philosophy*, 2007.

Ned Hall. Two mistakes about credence and chance. *Australasian Journal of Philosophy*, 82:93–111, 2004. Reprinted in Jackson and Priest 2004.

Joseph Y. Halpern. Lexicographic probability, conditional probability, and non-standard probability. In Johan von Bentham, editor, *Proceedings of the Eighth Conference on Theoretical Aspects of Rationality and Knowledge*, pages 17–20. Morgan Kaumann Publishers Inc. (San Francisco), 2001.

W. Harper, R. Stalnaker, and G. Pearce, editors. *Ifs.* Reidel, Dordrecht, 1981.

John Hawthorne. Craziness and metasemantics. *Philosophical Review*, 116:427–440, 2007.

Risto Hilpinen. Approximate truth and truthlikeness. In M. Przelecki, editor, *Formal Methods in the Philosophy of the Empirical Sciences.* Reidel, Dordrecht, 1976.

Carl Hoefer. The third way on objective probability: A skeptic's guide to objective chance. *Mind*, 116:549–596, 2007.

Frank Jackson and Graham Priest, editors. *Lewisian Themes. The Philosophy of David K. Lewis.* Oxford University Press, 2004.

Saul A. Kripke. *Naming and Necessity.* Harvard UP, 1980.

David Lewis. *Counterfactuals.* Oxford University Press, 1973.

David Lewis. How to Define Theoretical Terms. *Journal of Philosophy*, 67:427–426, 1970. Reprinted in Lewis 1983b (pp. 78-95).

David Lewis. New work for a theory of universals. *Australasian Journal of Philosophy*, 61:343–377, 1983a. Reprinted in Lewis 1999 (pp. 8-55).

David Lewis. Putnam's paradox. *Australasian Journal of Philosophy*, 62:221–236, 1984. Reprinted in Lewis 1999 (pp. 56-77).

David Lewis. Counterfactual dependence and time's arrow. *Noûs*, 13:455–76, 1979. Reprinted in Lewis 1986b.

David Lewis. Counterpart theory and quantified modal logic. *Journal of Philosophy*, 65:113–26, 1968. Reprinted in Lewis 1983b.

David Lewis. Humean supervenience debugged. *Mind*, 103:473–490, 1994. Reprinted in Lewis 1999.

David Lewis. *On the Plurality of Worlds*. Blackwell, 1986a.

David Lewis. *Philosophical Papers*, volume I. Oxford University Press, Oxford, 1983b.

David Lewis. *Philosophical Papers*, volume II. Oxford University Press, Oxford, 1986b.

David Lewis. *Papers in Metaphysics and Epistemology*. Cambridge University Press, Cambridge, 1999.

Barry Loewer. David Lewis's Humean theory of objective chance. *Philosophy of Science*, 71:1115–1125, 2004.

G.H. Merrill. The model-theoretic argument against realism. *Philosophy of Science 47*, 1980.

Daniel Nolan. Recombination Unbound. *Philosophical Studies 84.2-3*, pages 239–262, 1996.

Graham Oddie. Truthlikeness. *Stanford Encyclopedia of Philosophy*, 2007.

Karl R. Popper. *Conjectures and Refutations*. Routledge, London, 1963.

Hilary Putnam. Models and Reality. *Journal of Symbolic Logic*, 45:464–482, 1980.

Hilary Putnam. *Reason, Truth, and History*. CUP, 1981.

Gideon Rosen. Modal Fictionalism. *Mind*, XCIX:327–354, 1990.

Theodore Sider. *Four-Dimensionalism*. Oxford University Press, 2001.

Robert C. Stalnaker. *Ways a World Might Be. Metaphysical and Anti-Metaphysical Essays*. OUP, 2003.

Bas C. van Fraassen. A temporal framework for conditionals and chance. *Philosophical Review*, 89:91–108, 1980a. Reprinted in Harper et al. 1981.

Bas C. van Fraassen. *The Scientific Image*. OUP, 1980b.

Matthew Weiner and Nuel Belnap. How causal probabilities might fit into our objectively indeterministic world. *Synthèse*, 149:1–36, 2006.

J. Robert G. Williams. Eligibility and inscrutability. *The Philosophical Review*, 119:361–399, 2007.