

Probability Out Of Determinism

Michael Strevens

To appear in a slightly shorter version in *Probabilities in Physics*,
edited by Claus Beisbart and Stephan Hartmann, Oxford University Press.

ABSTRACT

This paper offers a metaphysics of physical probability in (or if you prefer, truth conditions for probabilistic claims about) deterministic systems based on an approach to the explanation of probabilistic patterns in deterministic systems called the method of arbitrary functions. Much of the appeal of the method is its promise to provide an account of physical probability on which probability assignments have the ability to support counterfactuals about frequencies. It is argued that the eponymous arbitrary functions are of little philosophical use, but that they can be substituted for facts about frequencies without losing the ability to provide counterfactual support. The result is an account of probability in deterministic systems that has a “propensity-like” look and feel, yet which requires no supplement to the standard modern empiricist tool kit of particular matters of fact and principles of physical dynamics.

ACKNOWLEDGMENTS

For helpful feedback, thanks to the editors, Marshall Abrams, Carl Hoefer, Aidan Lyon, Jacob Rosenthal, Chris Wüthrich, and participants in the Corridor Reading Group (Delia Fara, Boris Kment, Josh Knobe, Jill North, Jim Pryor, Brian Weatherson).

1. Deterministic Probability

A *deterministic probability* is a physical probability sincerely ascribed by some scientific theory to an outcome type that is produced by processes that are, deep down, deterministic or quasi-deterministic (meaning that all relevant fundamental-level probabilities are close to zero or one). Such a theory is imputing probabilities to a more or less deterministic world, then, and the question naturally arises—how can a deterministic world nurture stochasticity?

Deterministic probabilities are ascribed by statistical mechanics, both when it is applied in a classical context (as when it was first formulated) and when it is applied to small systems in a standard quantum mechanical context (since those systems will not normally undergo collapses in the short term). They are ascribed by population genetics and other evolutionary models in biology that allow for the possibility of drift (if the low-level biological processes underlying these phenomena are quasi-deterministic). They are apparently ascribed by a wide range of other models in biology and the behavioral and social sciences (though these models may vary in their “sincerity” concerning the probabilities’ observer-independent existence). Deterministic probability is, it seems, pervasive in the high-level sciences.

Some writers on probability have regarded deterministic probabilities as ersatz (Fetzer 1971; Giere 1973; Schaffer 2007). Even if they do not fit certain notions of “genuine probability”, however, that does not diminish the interest of the project of finding out what probabilistic theories such as statistical mechanics, evolutionary biology, and so on are telling us about the workings of the world. Put aside the question whether deterministic probabilities are “real probabilities” or not, then. The question I will ask in this paper is: what facts determine whether a deterministic probability ascription in physics or biology is true or false? What are the truth conditions for the probabilistic claims about deterministic systems made in so many of the sciences? That, I take it, is the problem of giving a metaphysics of deterministic probability.

There are two traditions in the interpretation of scientific probability ascriptions that have something substantive to say about deterministic probability. The first, de Finetti's (1964) subjectivism, understands the ascription of a deterministic or any other kind of probability as a projection of subjective probability onto the world. This approach has always been a minority view, perhaps because it simply raises the question "In virtue of what kinds of physical facts is it scientifically fruitful to project subjective probability in this way?", a question that looks rather like a rephrasing of the original metaphysical question—on my understanding of that question, at least.

The other interpretation of probability friendly to the deterministic variety is of course frequentism, which in its simplest versions equates probabilities with actual frequencies: the probability of a tossed coin's landing heads, for example, is equated with the frequency with which tossed coins actually land heads (or the limiting frequency, if there are infinitely many actual coin tosses). Few, perhaps, have endorsed this naive actual frequentism, but the idea that a scientific ascription of probability has its basis solely in facts about actual frequencies has attracted many thinkers (Reichenbach 1949; von Mises 1957; Salmon 1967), most recently under the guise of Humean "best system" accounts of lawhood (Lewis 1994; Loewer 2001; Cohen and Callender 2009) or something similar (Hoefer 2007).

In this paper I will investigate a third strategy for understanding deterministic probability, that of finding a basis for an outcome's probability in the properties of the physical dynamics that produce it. This strategy has a long and interesting history, first under the rubric of the "method of arbitrary functions" (von Plato 1983), and second, in foundational work on statistical mechanics, under the rubric of "modern ergodic theory" (Sklar 1993).

The dynamical approach, though attractive, will turn out to be not quite feasible; some kind of supplement to the dynamics is required. The natural place to look for the supplement is the actual frequencies: what I will ultimately advocate, then, is a metaphysics of deterministic probability largely

based in physical dynamics but with a frequentist element.

I will not attempt a complete exposition of the metaphysics of deterministic probability in this paper. Rather, I will focus on arguing that a metaphysics can be constructed that satisfies two principal desiderata. First, the building blocks of deterministic probability should not themselves be physical probabilities, or not always; thus, the metaphysics should allow for the existence of deterministic probabilities in a fundamentally deterministic world. Second, the inclusion of frequentist facts among the constituents of deterministic probability should not bring down upon the resulting metaphysics the most serious objections to frequentism, above all the objection to be laid out in section 2, that a frequentist metaphysics cannot distinguish between nomic and merely accidental statistical regularities. Other details of the account, however important, will be dubbed “secondary” for the purposes of this paper, and their consideration postponed.

2. Frequentism's Fundamental Flaw

Actual frequentism has been subjected to many powerful objections (Hájek 1996); the greatest of these is, in my view, that it cannot distinguish between accidental regularities and those that are to be accounted for by the nomic structure of world—a flaw that frequentism shares, unremarkably, with its close relative the naive regularity theory of deterministic laws (Strevens 2006).

This problem of accidental versus nomic regularities is the dramatic crux of all that is to follow, first because the dynamical approach to deterministic probability on which this paper is based appears to solve the problem, and second because the unavoidable supplementation of the dynamical approach with a frequentist component appears to unsolve the problem, rendering the dynamical approach no better than straight frequentism.

Imagine three universes. In the first is an urn of balls, half red and half white. Over the lifetime of the universe, 10,000 drawings are made from the

urn, approximately 50% of which produce a red ball. The second universe is identical, except that the 10,000 drawings are by some fluke 90% red. The third universe is richer and more culturally diverse. It contains a great and proud nation, the United States of Chimerica, that lasts for 800 years before succumbing to an influenza pandemic. Over that span of time, it has 200 presidents, 20 of whom have the first name James.

For the naive frequentist, the probability of an urn drawing producing a red ball in the first universe is 50%, the probability of the same in the second universe is 90%, and the probability of a US president being named James in the third universe is 10%. Such pronouncements miss crucial distinctions from both the scientist's and the metaphysician's point of view. There is something that is lawlike, or nomic, to the sequence of red balls drawn in the first universe: the drawings are produced by a nomic process, and they reflect the underlying nature of that process. The drawings in the second universe possess the first property but not the second. The presidential names lack even the first property. In both the latter two cases, we are inclined to say that the statistics—90% red, 10% James—are “mere accidents”.

This distinction between nomic and accidental statistical patterns matters deeply to us. If we were to discover that the statistical patterns studied by statistical physics, quantum physics, or population genetics were more like those in the second or third universes than those in the first—if we discovered that the patterns did not reflect the underlying nomic reality, or that there were no such underlying reality—our conception of the workings of the world would be revolutionized. Though the predictive utility of the theories that we had constructed on the basis of the patterns might not be diminished, their scientific status would be shattered.

I will not try to say here why the nomic/accidental divide matters so much to us. (Perhaps it has something to do with the desire for theories that not only summarize but explain the phenomena.) I will simply take for granted the importance of the distinction; it then follows that naive frequentism's

failing to respect the distinction—its positing probabilities wherever there are frequencies, regardless of whether the frequencies are nomic or accidental—is a reason for profound dissatisfaction. Many of the standard objections to frequentism have this dissatisfaction, I think, at their core.

The next step is to find a test that diagnoses the nomic/accidental regularity distinction and to use the test to frame an adequacy criterion for a metaphysics of deterministic probability.

The test will not surprise you: following a long tradition, I propose that a nomic frequency can be distinguished from an accidental frequency by its counterfactual robustness. Were 500 further drawings to have occurred in universe 1, or were the 10,000 drawings that did occur conducted under other circumstances (say, a minute later than in actuality), they would still have likely yielded about 50% red balls. But the same cannot be said of the 90% frequency in universe 2: another 500 drawings, or the actual drawings differently performed, would have more likely produced 50% than 90% red balls. Likewise, further elections in the USC, had they occurred, would not have been particularly likely to install 10% Jameses as president, either in the short or the long run.

This counterfactual test for a statistical pattern's being nomic rather than accidental may be transformed into a requirement on accounts of deterministic probability as follows. If the deterministic probability of a setup's producing an outcome of type e is p , then many or most counterfactuals of the following form should be true:

Had a further long series of trials been conducted, or had an actual long series been conducted under somewhat different circumstances, then the frequency of e events obtained would very likely be approximately equal to p .

The deterministic probabilities we know and love satisfy the robustness requirement. Of coin tosses, for example, it is correct to say (counting 500 tosses as a “long series”):

If I right now tossed this coin 500 times (in fact, I will not), I would likely get about one-half heads.

If this coin had featured Susan Sontag rather than Susan B. Anthony, then those 500 tosses would still have likely produced about one-half heads.

If the US produced a 2 cent coin, it would when tossed likely produce about one-half heads.

The more “scientific” deterministic probabilities—not least, those of statistical physics and population genetics—appear in the same sense to produce their corresponding frequencies robustly.

The requirement that probabilities generate frequencies robustly is not entirely well formed. For one thing, it does not specify an exact set of counterfactual antecedents relative to which robustness should be tested; this, however, including the question whether some antecedents should concern short series of trials, is a matter to be finessed once the metaphysics is complete.

Another worry is more urgent: the conditionals in question use terms—*typically*, *likely*—that themselves require some interpretation. I see two ways to understand such terms, or more exactly, I see two distinct kinds of counterfactual conditionals that can be expressed with the “likely” formulation; call them type 1 and type 2 conditionals.

Type 1 conditionals concern counterfactual probability distributions: in a Stalnaker/Lewis semantics, the type 1 conditional *If A then likely B* expresses the fact that in all the closest worlds where *A* holds, the physical probability of *B* is high. The kind of robustness we want in the frequencies produced by genuine deterministic probabilities is not, however, correctly diagnosed by type 1 conditionals. To see this: Suppose that in universe 2 above, I had performed 500 further drawings. However these drawings come out, the frequency of red in the resulting universe will be about 90% (because the results of the 500 drawings are swamped by the results of the

10,000 extant drawings). It follows that, were naive actual frequentism to be correct, the relevant type 1 conditional would be true: were I to perform 500 further drawings, the probability of red would still be, according to actual frequentism, about 90%, and therefore the probability of those 500 drawings producing about 90% red balls would be very high. In short, actual frequentism performs better on the robustness test than it ought to (though not perfectly) if the relevant conditionals are interpreted as being of type 1. The problem is that type 1 conditionals themselves contain a reference to physical probability. An unscrupulous account of physical probability can use this reference to rig the test in its own favor, as frequentism to some extent does.

Type 2 conditionals promise to be more discerning. The type 2 conditional *If A then likely B* is true just in case in nearly all the closest worlds where A holds, B also holds.¹ As you can see, actual frequentism will fail the test for robustness in universe 2 if the test's conditionals are type 2: it is not the case that, were I to conduct 500 further drawings, I would in most closest worlds draw about 90% red balls. Thus, the type 2 interpretation comes much closer than the type 1 interpretation to capturing the nomic/accidental distinction.

To establish the legitimacy of type 2 conditionals requires further work.

1. A similar semantics for “likely” conditionals is advocated by Bennett (2003, §97). But Bennett holds, following Bigelow (1976), that such conditionals are simply one example of a kind of conditional that, like type 1 conditionals, assigns a probability to the consequent: for Bennett (and presumably Bigelow), to say “If A then likely B” is to say “If A had occurred then B would have had a high probability”, except that the probability in question is not physical probability, as in a type 1 conditional, but some kind of logical probability. (At least, Bigelow says that it is logical. I am not sure what Bennett thinks it is; if he thinks it is physical, then he is simply assimilating type 2 to type 1 conditionals and giving a theory of physical probability along the way.)

In my view, type 2 conditionals are not about probability at all—at least, not about any of the kinds of probability recognized by philosophers: subjective, physical, and epistemic (logical probability being a sub-type of epistemic probability). There is no interpretation of probability to be found in my semantics of type 2 conditionals, then, only an interpretation of a certain variety of “likely” talk, a kind of talk that makes sense, I will argue in section 4.3, even when there are no relevant physical or epistemic probabilities in the vicinity.

First, some sense must be given to “nearly all”; it seems wise to allow that the threshold value is both vague and to some degree sensitive to context (though not so variable that the characterization of what it is to produce frequencies robustly is entirely meaningless). Second, some kind of measure must be provided over the nearby worlds with respect to which to evaluate the proportion of *B*-worlds among closest *A*-worlds. Finally, if the metaphysics I will give for deterministic probability is to be genuinely reductive, it must be established that the measure in question is not a physical probability measure, or at least that it is not always a physical probability measure. These concerns will be addressed later in the paper; for now, I will simply assume that the correct test for a nomic or robust frequency, and thus the correct adequacy condition for an account of deterministic probability, is framed in terms of conditionals of type 2.

Naive frequentism disappoints because it does not satisfy the robustness criterion. Other versions of frequentism do better. Hypothetical frequentism more or less takes the criterion as the definition of physical probability, though typically with a reference to limiting frequencies in infinitely long series of trials that raises various difficulties (Hájek 2009). Analyses inspired by the best system account of laws can by various means avoid ascribing probabilities where frequencies are intuitively accidental, though how they should deal with universe 2 above is unclear (Hoefer 2007).

I will not investigate these sophistications of the frequency account in this paper, however. They do not inspire me, for the following reason: they identify deterministic probabilities with robust frequencies, where it would make more sense to identify the probabilities with the physical matters of fact that *explain* robustness. For example, what makes the probability of heads on a coin toss one-half ought not to be the fact that the frequency of heads is robustly one-half, but the facts about the world—the physics of coin tosses and so on—that explain this robustness. Hence my interest in dynamical accounts of deterministic probability such as the method of

arbitrary functions.

3. The Method of Arbitrary Functions

A metaphysical reduction of deterministic probability must derive probability distributions from something non-probabilistic. The method of arbitrary functions has the appearance of conjuring probability out of physical dynamics without appealing either to frequencies or to any other probabilistic fact. This appearance does not survive even cursory philosophical scrutiny: the method does not create probabilities *ex nihilo*, but rather requires some probabilistic material to work with. It asks so little of that material, however, that it seems that frequencies might just be capable of supplying what is needed: they may not be able to bear the full weight of a metaphysics of deterministic probability, but they can support one corner of the edifice. That, at least, is what I will argue.

Ever since probability found itself, in the mid-nineteenth century, reified in physical theory, the thought that such probabilities could have their basis in a certain physical dynamics has attracted philosophical writers. What has come to be called the method of arbitrary functions has been advocated by von Kries, Poincaré, Reichenbach, Hopf, and a number of others (von Plato 1983). In my view, it finds its most mature form in Hopf's presentation (Hopf 1934); it is Hopf's line of thinking that I will follow.

Consider a simple wheel of fortune, that is, a wheel painted in equal numbers of equal-sized red and black sections like a roulette wheel (without the zeroes). The wheel is spun on its central axis and allowed to come to rest, with a fixed pointer then indicating the outcome, *red* or *black*.

The probability of *red* on such a device is evidently one-half, a gut feeling easily confirmed by statistical testing. But the physical process producing *red* is deterministic, or as close as quantum mechanics allows (which is very close). The one-half probability is therefore a deterministic probability. What

is its basis?

If there were a determinate probability distribution over the initial conditions of spins on the wheel, then the basis would presumably be that distribution plus the aspects of the wheel's physical dynamics that determine which initial conditions result in *red* and which in *black*. More exactly, the probability of *red* would be the probability assigned by the initial condition distribution to the set of initial conditions resulting in *red*.

But that, of course, is simply to move the locus of, rather than to solve, the problem of deterministic probability: where does the probability distribution over initial conditions come from? Perhaps the process by which initial conditions are generated—the features of croupiers' physiology in virtue of which the wheel of fortune acquires its initial speed in any spin of the wheel—is at bottom stochastic rather than deterministic. Then you would have your initial condition distribution, but you would have said nothing useful about deterministic probability. Suppose that the croupier is as deterministic as the wheel. What, then, is the ultimate source of the probability of *red*?

Hopf and others thought that a good part of the answer might come from the physical process that converts initial conditions into final outcomes, that is, that takes an initial spin speed and determines the final outcome, *red* or *black*. (In what follows, I assume for simplicity's sake that the outcome is entirely determined by the wheel's initial spin speed.)

Take a look at the shape of the dynamics of that process for a typical wheel of fortune, as represented by a function that maps *black*-producing initial spin speeds to zero and *red*-producing initial spin speeds to one—what I call the wheel's *evolution function* (figure 1).

You will observe that the function has a certain characteristic structure: the ratio of the “gray” to the “white” sections (representing spatiotemporally contiguous sets of initial speeds producing *red* and *black* respectively) in any fairly small neighborhood is the same, 1 : 1. To put it another way, in any small but not-too-small range of spin speeds, the proportion of speeds

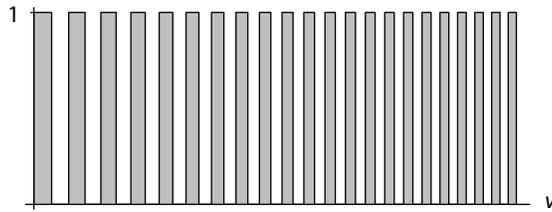


Figure 1: Evolution function for a simple wheel of fortune, mapping initial spin speed v to either 1 (*red*) or 0 (*black*). The area under the function, corresponding to *red*-yielding values of v , is for clarity's sake shaded gray.

producing *red* is one-half. I call this property of an evolution function—that within any small but not too small neighborhood, the proportion of initial conditions producing a given outcome is the same—*microconstancy*. I call the proportion in question the evolution function's *strike ratio* with respect to that outcome.²

Putting these terms to use: the dynamics of the wheel is microconstant (relative to the outcomes *red* and *black*) with a strike ratio for *red* of one-half. Had the wheel been painted so that one-third of its colored sections were red and two-thirds were black, its dynamics would still have been microconstant, but with a strike ratio for *red* of one-third. Had the wheel been constructed so that it wobbled as it rotated around its axis, it would (depending on the physical details) quite possibly not have been microconstant. The fact of microconstancy and the value of the strike ratio are therefore determined by the dynamics of the wheel. (The notions of microconstancy and of a strike ratio, along with the other technical aspects of the techniques described in this paper, are laid out more carefully and with greater formality in Strevens (2003), as is the extension of the notions to systems with more than one initial condition.)

2. Note that microconstancy is relative not only to an outcome but also to a way of measuring the initial conditions. Later uses of the notion of microconstancy ought therefore to allow for measure-relativity; this is taken care of in notes 8 and 18.

It is of course no coincidence that the strike ratio for *red* is equal to the probability of *red*. But how can you know this, if the probability of *red* is determined in part by the probability distribution over initial spin speed, and you have very little idea what that distribution looks like? Because you can see, intuitively, that almost any probability distribution over spin speed will induce a probability for *red* approximately equal to the strike ratio for *red*. Figure 2 illustrates this physico-mathematical intuition: for two rather different speed distributions, the probability of *red*, equal to the probability assigned by a probability distribution to the “gray” areas—that is, the proportion of the area under the probability density that is shaded gray—is the same, equal to the strike ratio of one-half. You might say that the probability is in these cases in effect determined by the strike ratio alone.

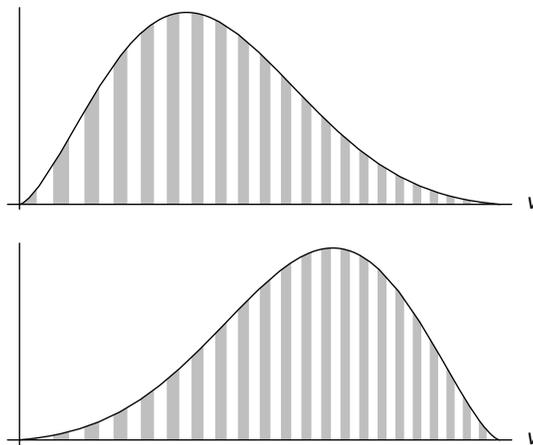


Figure 2: Different spin speed distributions induce the same probability for *red*, equal to the strike ratio for *red* of one-half

It is this observation that catches the attention of the would-be metaphysician of deterministic probability. If the probability is determined by the strike ratio alone, and the strike ratio is determined by the deterministic dynamics of the device in question, then the probability is determined by deterministic facts alone. A theory of deterministic probability falls out of

the mathematics of Hopf et al. and into the lap of the ontologist more or less unbidden.³ Further, because the theory identifies deterministic physical law as the sole determinant of deterministic probability, it easily satisfies my principal desiderata for a metaphysics of probability: it makes no use of probabilistic building blocks, and it seems sure to offer a conception of probability that, like physical laws themselves, generates frequencies robustly, so satisfying the counterfactual requirement articulated in section 2.

If only the phrase “determined by the strike ratio alone” were not an exaggeration. If only it were true that any probability distribution whatsoever over the wheel of fortune’s initial spin speeds would induce a probability equal to the strike ratio—then you could define the probability of *red* as nothing over and above, but something exactly equal to, the strike ratio for *red*.⁴

It is not to be. Examine figure 3. Here is a spin speed distribution that induces a probability for *red* that is significantly greater than one-half (well over one-half the area under the density function is shaded gray). As you can see by comparing the figures, what the distributions in figure 2 that induce a probability equal to the strike ratio have, and what the distribution in figure 3 that fails to do so lacks, is a certain kind of smoothness: the distributions in figure 2 are roughly flat over any small interval. I call this property of near-flatness over small intervals *macroperiodicity* (where what it takes to be

3. In understanding the method of arbitrary functions as offering a full-blown metaphysics of probability, I am going beyond my texts. There are hints of this picture in Poincaré and Hopf, but I suspect that they would resist the view in its starkest form. Von Kries, by contrast, embraces it, but with certain qualifications and complications that set his thinking apart from the simple story examined here. (Rosenthal (2010) both explains and advocates an improved version of von Kries’ ideas.) The aim of what follows, then, is to develop a certain line of thought, rather than to follow the twists and turns of a particular intellectual biography.

4. Even this thinking is somewhat wishful. Any perspicuous justification for such a definition must assume that *some* probability distribution exists over initial conditions, and this is eminently disputable in the case where the initial conditions are deterministically produced. But there is no need to dwell on this point, as it is about to be superseded.

“small” is determined by the scale of the evolution function’s microconstancy).

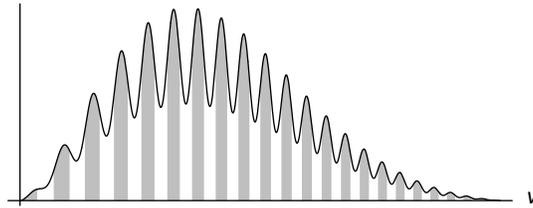


Figure 3: Not every initial spin speed distribution induces a probability for red of one-half

To derive a probability from a strike ratio, then, a further assumption about the initial condition distribution is required, the most obvious such assumption being that the distribution is macroperiodic.⁵ And this assumption in turn requires the existence of an initial condition probability distribution.

Hopf and others point out, in an attempt at amelioration, that as the average spin speed on the wheel of fortune increases, the gray and white bars in the evolution function alternate more and more quickly, and as a result the requirement of macroperiodicity becomes weaker and weaker, since the regions over which approximate uniformity is required become smaller and smaller. In the limit, the macroperiodicity requirement appears to vanish altogether: all that is required for the initial condition distribution to induce probabilities equal to strike ratios is that the distribution have a probability density, or in the jargon, that it have the property of absolute continuity. Choose whatever function you like as the initial condition density, then—choose an *arbitrary function*—and you will get a probability equal to the relevant strike ratio.

This is an interesting observation, but its relevance to the probability of

5. Macroperiodicity is a sufficient rather than a necessary and sufficient condition: some non-smooth densities will induce a probability equal to the strike ratio “by chance”, but these flukes are, I suspect, of little physical interest.

red on a real wheel of fortune is, to put it delicately, unclear. If the wheel is operated by a normal human, its spin speeds will not exceed a certain threshold. How, then, could the properties of the wheel when spun at speeds beyond that threshold—let alone its limiting properties as the speed goes to infinity—be relevant to what actually happens? Surely (Strevens 2003, §2.A) it is the properties at actual, not counterfactual, speeds that matter.⁶

The method of arbitrary functions almost works. When you have a device with an evolution function that is microconstant relative to some outcome, the probability for that outcome induced by the initial condition distribution will be equal to the strike ratio—a quantity determined by the device’s dynamics alone—provided that a rather weak constraint on the initial condition distribution holds. The probability is *almost* determined by dynamics alone. But *almost* is not worth much in the world of absolutes where metaphysics makes its home. Thus, though the urge to do something metaphysical with microconstancy is easily excited, the philosophical means seem to be lacking.

4. Microconstant Probability

4.1 *A New Hope*

I assumed in the previous section that the probability of *red* on the wheel of fortune is whatever probability is induced by the initial spin speed distribution and the wheel’s dynamics, or in other words, that the probability of *red* is determined by the probability that the initial condition distribution assigns to the set of *red*-producing initial spin speeds.

6. Even if the behavior of the wheel in the limit were of immediate practical relevance, the determination of the probability of *red* from the strike ratio for *red* requires that the initial condition distribution have the property of absolute continuity (a property that some probability distributions lack). If there is to be a fact of the matter about its properties, the distribution must presumably exist. So we are back at square one, facing the problem of where a deterministic probability distribution comes from. Compare note 4.

All subsequent problems may be traced back to this assumption, since it more or less ensured that the conditions under which the probability of *red* could be set equal to the strike ratio for *red* included a constraint on (or at least, the existence of) a probability distribution over initial spin speed. And so the problem of finding a deterministic basis for probabilities was not solved but merely pushed one step back along the causal chain.

I will try something new: I will sever the link between the initial condition distribution and the probability of *red*. This means that I will no longer assume that there is a probability distribution over spin speed. It also means—more radically—that even if there is an initial condition distribution over spin speed, it will not play a part in defining the probability of *red*.

In the previous section, the implicit metaphysical equation was something like this:

Deterministic probability equals (or is determined by) properties of the physical dynamics, in particular microconstancy, plus properties of the physical probability distribution over initial conditions, in particular macroperiodicity.

A metaphysics of deterministic probability along these lines is suggested by, among others, Sober (2010), for whom the probability of heads on a tossed coin is derived from facts about a physical probability distribution over the toss's initial speed and loft.

The new equation I propose to put in place of the old is:

Deterministic probability equals (or is determined by) properties of the physical dynamics, in particular microconstancy, plus properties of the frequencies of initial conditions.

Where there was the probability distribution over initial conditions there are now facts about the initial condition frequencies. But—and this is crucial—I am not giving a frequentist account of the initial condition distribution. That would be to accept the first equation above and then to supplement it with a

frequentist interpretation of the initial condition probabilities. I am rejecting the first equation. A probability distribution over initial conditions, whether it exists or not, has no place in my metaphysics of deterministic probability. The role played by such a distribution in the first equation is played by non-probabilistic facts in my account—non-probabilistic facts about frequencies. Consequently, my metaphysics of deterministic probability assigns physical probabilities only to the outcomes of microconstant processes. It neither requires nor implies the existence of physical probabilities attached to those processes' initial conditions.

This does not rule out the existence of another kind of physical probability that is induced by a physical initial condition probability distribution. I call these other probabilities *derived physical probabilities*. The derived physical probability of a given process's producing an outcome, then, is the probability assigned by the probability distribution over the process's initial conditions (should such a distribution exist) to the set of initial conditions that produces the outcome. The probability distribution over initial conditions may be any kind of physical probability: it might be the kind of deterministic probability defined in this paper; it might be an irreducible probability distribution (from quantum physics or wherever); it might be some other kind of deterministic probability; it might be a mix of these (if there is a quantum probability distribution over one kind of initial condition and a deterministic probability distribution over another, for example); and it might of course itself be a derived physical probability.⁷

If there is a physical probability distribution over the initial conditions of a microconstant process, it follows that the outcomes of the process may fall under two probability distributions: a derived physical probability distribution and a distribution of the kind of free-standing deterministic probability to be defined shortly. I will show in section 4.4 that the two probability

7. The networks of probabilistic dependence implied by derived probability are discussed in Strevens (2003), §2.4.

distributions will live in harmony: where both distributions exist, they will agree in all their probability ascriptions.

4.2 *Deterministic Probability Provisionally Defined*

My initial attempt at a metaphysics of underived deterministic probability, implemented for the wheel of fortune, is as follows. (Two alternative ways to use microconstancy or something similar to define deterministic probability are proposed by Abrams (2000) and Rosenthal (2010). Like the provisional definition set out in this section, Abrams' definition makes use of constraints on finite subsets of actual initial conditions; Rosenthal's definition was characterized in note 3.)

Say that a set of actual wheel of fortune spin speeds is *macroperiodically distributed* if the probability density that best summarizes the frequencies of speeds in the set is macroperiodic (more on this shortly). Then stipulate that the outcome *red* on a spin of a wheel of fortune has a deterministic probability if:

1. The dynamics of the wheel is microconstant with respect to *red*, and
2. The actual initial conditions of nearly all long series of spins on wheels of fortune form macroperiodically distributed sets.

The deterministic probability, if it exists, is stipulated to have a value equal to the strike ratio for *red*.⁸

Generalizing, the event of a system *S*'s producing an outcome of type *e*—the event of a single trial on *S*'s producing *e*—has a deterministic probability equal to *p* if

8. More exactly, for the probability to exist, there must exist a measure on the initial conditions relative to which the dynamics and initial conditions are respectively microconstant and macroperiodic (see note 2). This is a weaker requirement than you might think (Strevens 2003, 86–87); it will be strengthened in section 4.3.

1. The dynamics of S is microconstant with respect to ϵ , and has strike ratio p , and
2. The actual initial conditions of nearly all long series of trials on systems of the same type as S form macroperiodically distributed sets.

The definition gives you a probability for a binary outcome space, but you can generalize to, say, a probability distribution over a real-valued quantity if the evolution function for every Borel set is microconstant.

Call the kind of deterministic probability so defined *microconstant probability*. As you can see, a microconstant probability is determined by, or if you like constituted by, two kinds of properties: the properties of the system's dynamics in virtue of which it is microconstant with a certain strike ratio, and the property of the macroperiodicity of the system's and all similar systems' actual initial conditions in most long series of trials.

The definition of the macroperiodicity of a system's actual initial conditions raises three questions. First, what criteria are used to determine what count as systems of the same type? Second, what exactly does it mean to say that a set of actual conditions is macroperiodically distributed? Third, what counts as a "long series of trials"? (There is no need at this stage to precisify "nearly all"; let me just say that it is not a technical term.)

I will postpone the first question to section 4.4. The second question may be answered as follows. To determine whether some large set of actual values of a variable ν have a macroperiodic distribution of frequencies, divide up the range of ν into small, equally sized intervals.⁹ Then calculate the frequency with which values of ν fall into each of these intervals (if there are infinitely many actual values of ν , take the limiting frequencies). The results may be

9. What counts as "small" should be relativized in the obvious way to the "smallness" of the relevant evolution function's microconstancy—very roughly, you want intervals that are of the same magnitude as the "gray" and "white" bars in the evolution function. (Strevens (2003, §§2.33, 2.72) gives a more precise formulation.) Likewise, the measure used in assessing equality of size should be the same measure with respect to which the evolution function is microconstant (see note 2).

summarized by a plot of points; connect the dots. The resulting line is the “density” of the frequency distribution. The distribution is macroperiodic if its density is approximately macroperiodic—if it does not deviate too often or too much from strict macroperiodicity (the question how to precisify this requirement is secondary; it will not be pursued here).

What about the third question—the question of what counts as a “long series of trials”? “Long” must be long enough that the initial conditions stand a mathematical chance of determining a macroperiodic density. Further, the series should have some kind of spatiotemporal contiguity. The series of trials need not, however, be conducted all at once by a single person. They might be the spins on wheels of fortune made by a single person over a lifetime. Or they might be the spins made on a single wheel of fortune by a variety of people over the course of a week. That leaves considerable latitude in determining which are the long series of trials, but I do not think that the exact nature of the rules for individuating long series will make much of a difference in practice; in a world like ours, the same microconstant probabilities will exist no matter how exactly you carve up trials into sequences of trials, provided that the sequences are long and in some sense contiguous.¹⁰ This remark will be amplified in section 4.3, where it will also become clear why I go by way of “long series”, rather than simply requiring for the existence of microconstant probability that the complete set of initial conditions of all actual spins be macroperiodic.

The definition of deterministic probability is not arbitrary: the constituents of a deterministic probability so defined both predict and explain the patterns of outcomes produced by the devices to which they are attached. For reasons of space, I will have to provide you with only the rudiments of the reasons why, referring you to other research for the details.

10. Even with this liberal conception of what constitutes a long series, could there be large numbers of trials, perhaps even a majority, that belong to no such series? If so, you might want to add to the definition of microconstant probability a further requirement that these “stray” trials form a macroperiodic set.

First, prediction. It is easy to see that the microconstant probability of *red* will predict the fact that the frequency of *red* outcomes on a “long series” of actual spins of a wheel of fortune is one-half. How, though, does it predict a one-half frequency on a subsequence of such a series? How does it predict, for example, that some particular subsequence of 100 spins, within a much longer sequence of spins, will yield about 50 *red* outcomes? (Assume for the sake of the argument that the subsequence does not itself qualify as a “long series”.)

You might think that, in order to predict the one-half frequency of *red* in the 100 spins from the macroperiodicity of the spin speeds of the longer sequence, you would have to assume that the speeds of the 100 spins were drawn randomly from the set of all spins in the sequence. But no such assumption is needed (and just as well). Because of the microconstancy of the wheel of fortune, a much weaker assumption will suffice, namely, the assumption that the density representing the frequencies with which sets of 100 spin speeds of various magnitudes appear together—the joint density over the initial conditions of actual 100-spin sequences—is macroperiodic.¹¹ This same assumption is sufficient for the stochastic independence of the 100 outcomes. Thus, a microconstant dynamics can take a non-independent distribution of initial conditions and produce an independent distribution over outcomes, which is enough to derive all the predictions you need about subsequences of long series of trials.

An example will illustrate the form of the general independence result. Consider any two sequential trials on a wheel of fortune. Are the outcomes of such trials independent? Because of the wheel’s microconstant dynamics, the answer is affirmative if the joint distribution over sequential pairs of initial spin speeds is macroperiodic. The form of the joint distribution is determined by the relevant frequencies: you take as your data points all the

11. The “random draw” assumption is also an assumption about a joint density, but an extremely strong one.

pairs of initial conditions for actual pairs of sequential trials, and let these points determine the shape of a density. (The points and the density will inhabit a space with twice as many dimensions as the space for representing the initial conditions of a single trial.) Independence holds if this frequency-based density is sufficiently smooth. The same goes for groups of any number of outcomes. Is the outcome of a spin on the wheel of fortune independent of the outcomes of the last spin but one and the spin before that? Construct a joint density representing the distribution of all actual sets of three initial spin speeds that bear the relevant relation to one another, that is, sets that comprise the speeds of the n^{th} , $(n - 2)^{\text{th}}$, and $(n - 3)^{\text{th}}$ spins in a series. If the density is macroperiodic, independence holds. (Further details are given in Strevens (2003), chapter three.)

I have not given you quite what I promised. The microconstant probability of *red* was supposed to predict the one-half frequency of *red* in 100 spins, but on my story, it is the microconstant probability together with the additional property required for independence (the macroperiodicity of a certain joint distribution) that yield the prediction. That, however, is exactly as it should be: as a matter of general probabilistic fact, a probability of one-half is not on its own enough to predict a frequency of one-half in a finite set of outcomes (even with high probability); independence, or something like it, is needed as well. Indeed, facts about independence are typically built into the claim that a certain phenomenon conforms to a certain probability distribution. To say that coin tosses have a Bernoulli distribution with probability one-half is, for example, not only to say that the probability of heads is one-half, but that individual tosses are stochastically independent.

Next, and only briefly, explanation. That a microconstant probability predicts frequencies from facts about the initial conditions and causal dynamics that produce those frequencies goes a long way toward establishing that microconstant probability has explanatory power. A number of interesting problems remain, however. How (if at all) do microconstant probabilities

explain the properties of very short series of trials? Of single outcomes? Can a microconstant probability explain an atypical frequency? Do microconstant probabilities, because they are constituted by facts about all initial conditions and not just those directly causally relevant to a particular series of trials, contain information that is explanatorily irrelevant to a particular series? The answers to these questions and others are discussed in chapters ten and eleven of Strevens (2009).

4.3 *Robustness Recovered*

Facts about frequencies give microconstancy—the dynamical property at the heart of the method of arbitrary functions—a grip on the actual world it would otherwise lack, and thus allow an account of deterministic probability that turns on physical dynamics. But the inclusion of these same facts saps dynamical probability of its modal richness: there is not enough to microconstant probability, as provisionally defined in the previous section, to ensure that the frequencies generated by the dynamics are produced robustly.

For example, the one-half microconstant probability for heads is founded in the microconstancy of the tossing process and macroperiodicity of the actual spin speeds of the actual worlds' tossed coins. These facts are not quite sufficient in themselves to ground the counterfactuals that diagnose robustness, such as

If I right now tossed this coin 500 times, I would likely get about one-half heads.

My plan in this section is to add to the definition of microconstant probability whatever provides the basis for counterfactual truths of this sort, ensuring that deterministic probability and the kind of robustness characteristic of nomic as opposed to merely accidental statistical patterns go hand in hand.

I will focus on the conditional in the previous paragraph. Why is it true? In part because the evolution function of the coin is microconstant with strike

ratio one-half, and in part because of the truth of a further counterfactual (assuming that the outcome of a coin toss is determined by a single initial condition, “spin speed”):

If I right now tossed this coin 500 times, the initial spin speeds of my tosses would likely have a macroperiodic distribution.¹²

Recall that I interpret these robustness-diagnosing counterfactuals as being of type 2: the conditional above is not saying that the distribution has a high physical probability of being macroperiodic in all relevantly close possible worlds, but that the distribution is macroperiodic in nearly all such worlds.

The question, then, is why this macroperiodicity counterfactual is true and, more generally, why the macroperiodicity of long series of actual initial spin speeds is a robust property (I am assuming as before that 500 tosses constitute a “long series”). Macroperiodicity’s robustness has two aspects. First, the actual initial spin speeds of long sequences of tosses not only form a macroperiodic set, but would likely continue to do so were the tosses conducted under a wide range of counterfactual suppositions. Second, the actual initial spin speeds of long sequences of tosses that might have been made but were not would also likely be macroperiodically distributed, were the trials in question somehow to occur. Note that frequencies of initial spin speeds are not in general robust: the statistical profile of the spin speeds I generate may, and presumably will, change as my physiological condition changes. To say that macroperiodicity is robust is to say that, however the profile varies, it always has the property of macroperiodicity. It is this stable element of the initial condition frequencies that underwrites the robustness of the probability of heads.

I will explore macroperiodicity’s robustness by unraveling—in a schematic way—the reasons that large sets of actual spin speeds are in fact

12. Note that a large number of tosses are needed both for the first counterfactual to seem true and for there to be sufficiently many initial conditions to define a rough density, and so for the second counterfactual to have any hope of being true.

macroperiodic. The same facts that explain actual macroperiodicity will turn out to explain the robustness of macroperiodicity.

I am about to toss a coin 500 times. (These are therefore actual, not counterfactual, tosses.) Call my state at this time, along with the state of my environment, the pre-toss state. Some set of parameters causes me to toss at the particular moments that I do. Perhaps I am handed a tossing schedule, or perhaps the timing of the tosses is prescribed by the pre-toss state itself. Either way, the spin speeds that I impart to the coins are determined by three things:

1. The pre-toss state,
2. The facts about the timing of the tosses (possibly determined by the pre-toss state), and
3. The physics of the coin (e.g., its weight and size).

Why are most such series of 500 spin speeds distributed macroperiodically? A minimal answer: the typical actual pre-toss state, together with the typical actual set of timings and the typical actual coin, usually produce—because of such and such properties of the physiology of tossing—a macroperiodic set of spin speeds. But a richer answer is possible.

First, it is unnecessary to restrict the claim to actual coins. Coins that will never exist, such as the Susan Sontag dollar, would also, given the typical pre-toss state and timings, usually produce a macroperiodic set of spin speeds (assuming of course that they are of roughly normal size and shape). Given the current state of my physiological knowledge, I cannot prove this, but I think that there is extremely strong empirical evidence in its favor, namely the macroperiodicity of the spin speeds produced when tossing a wide range of physically different coins.

Second, and by much the same reasoning, it is unnecessary to restrict the claim to actual sets of timings. Just about any feasible set of timings would, in conjunction with a typical actual pre-toss state and a normal coin, produce

a macroperiodic set of spin speeds. (It may help to imagine the timings as exogenously provided, that is, as coming from the outside rather than being prescribed by the pre-toss state.) In other words, nearly any actual pre-toss state, in conjunction with nearly any feasible set of timings— not just actual sets of timings— and with any normal coin, will produce a macroperiodic set of spin speeds.

This is almost enough to explain not only the actual preponderance of macroperiodicity but also its robustness. One more thing is needed. So far, by an “actual pre-toss state” I have meant the state that precedes an actual series of tosses. But many of the robustness-diagnosing conditionals concern times when a series of tosses could have been made but was not. I therefore expand my definition of *actual pre-toss state* to actual states at moments when a series of tosses could realistically have been made, regardless of whether the series was actualized. Note that, as the words imply, all such states are actual; I am not here making any claims about possible but non-actual pre-toss states.¹³

We have good reason, I suggest, to believe that there is no relevant difference between pre-toss states that are followed by actual tosses and those that are not. It follows that we have good reason to believe that my previous claim about pre-toss states remains true under the new interpretation: nearly any actual pre-toss state (now meaning any state that *might* have set the stage for a series of tosses), in conjunction with nearly any feasible set of timings and with any normal coin, will produce a macroperiodic set of spin speeds. That is enough to underwrite and explain the truth of the type 2 counterfactuals, and thus the robustness of macroperiodicity.

Let me explain why, using the Lewis/Stalnaker approach to the semantics of counterfactual conditionals as interpreted by Bennett (2003). Lewis held

13. It is tempting to pursue the investigation in this direction, but as you will see, it is unnecessary. This is a good thing, because to characterize and quantify the range of possible pre-toss states would be a formidable challenge. How wide a range of possibilities should we consider? Should we consider different possible evolutionary histories for the human race, resulting in different physiologies?

that the conditional *If A had occurred at time t, then B* is true if in all the closest possible worlds in which *A* occurs at *t*, *B* comes to hold.¹⁴ I have endorsed a natural extension of Lewis's truth conditions to conditionals of the form *If A had occurred at time t, then B would likely have occurred*, true on a type 2 interpretation if *B* holds in *nearly all* the closest possible worlds in which *A* occurs at *t*. The semantics is then filled out by giving a characterization of the relevant sense of closeness.

A world in which *A* occurs at *t* is close (according to Bennett) if it satisfies the following conditions:

1. Its state until shortly before *t* is identical to the state of the actual world,
2. Shortly before *t*, a conservative deviation from the actual world's history takes place, leading to the occurrence of *A* at *t*, and
3. After *t*, events unfold as dictated by the actual laws of nature.

The closest worlds are those in which the deviation in the second stage is most conservative. I will not go into the quantification of closeness here; suffice it to say that a conservative deviation involves (in the deterministic case) only a small and local violation of the actual laws of nature to set it off—a “small miracle”, in Lewis's language—and has very few side-effects. In other words, it changes as little as possible about the actual world, except insofar as it brings about *A* and the consequences of *A*.

Consider the conditional at hand, then: If I tossed this coin 500 times, the initial spin speeds of my tosses would have a roughly macroperiodic distribution. The worlds relevant to assessing the conditional—the closest worlds in which I conduct the tosses—will be those that satisfy the following criteria.

14. I will not worry here about my tacit use of what Lewis calls the “Limit Assumption”. While I am excusing my assumptions, I should also note that the Bennett semantics endorsed here does not apply to all counterfactuals; however, it does apply to the sorts of counterfactuals that capture the modal vigor of deterministic (and other) physical probabilities.

1. They are identical to the actual world until shortly before the series of tosses hypothetically takes place. It follows that in these worlds, I and my environs will be in a state identical to the actual pre-toss state.
2. Shortly before the hypothetical tosses, something happens that motivates me to toss the coin 500 times, without affecting anything else that is relevant to the tosses, as far as possible. (Perhaps someone appears and pays me \$100 to toss the coin 500 times, or perhaps instead of idly playing with my hair I begin to idly play with a coin.)
3. Otherwise, everything obeys the actual laws of nature.

What happens in such worlds is predicted by my claim above, that for nearly all actual pre-toss states, and for nearly any feasible schedule of tosses, the initial speeds of the tosses will be macroperiodically distributed. The closest worlds in which the tosses take place are therefore nearly all worlds with macroperiodic sets of initial speeds. The counterfactual conditional is true, then, as for similar reasons are the many other conditionals that diagnose the robustness of macroperiodicity. Thus, the processes to which microconstant probabilities are attached typically produce the right frequencies not only in actuality but in counterfactuality, as desired.

Let me reconsider the worries about type 2 “likely” counterfactuals mentioned in section 2. First, such counterfactuals invoke a threshold. What proportion of worlds is “nearly all”? In the present context, that depends on what you think it would take to make a frequency adequately robust. Let me suggest 95% as a reasonable choice; you may prefer a different value.

Second, such counterfactuals invoke a measure of possible worlds. It would be spectacular to provide a theory of the appropriate measure for every type 2 counterfactual. Here I aim for something more modest, a measure for those conditionals about macroperiodicity that are immediately relevant to my inquiry. The problem can be divided into two. You need a measure over different timing schedules for tosses; since these are finite sets of real

numbers, a version of the standard measure will do nicely.¹⁵ And you need a measure over different “most conservative” ways of altering the actual world to implement such a schedule. These “ways” will be physical processes; again I suggest the standard measure over the usual physical quantities: time, distance, and so on. The justification for these choices of measure is that they capture the actual truth conditions for type 2 conditionals, as determined (I presume) by us and our language. As far as I can see, the conventional measures I have suggested match our thinking fairly well.¹⁶

Third, in applying the standard measure as suggested, we do not think of it as a physical probability distribution, that is, we do not think of the standard measure as giving us a probability distribution over tossing schedules or over possible deviations from actuality leading to the application of such schedules. How do I know? Because I am morally certain that the macroperiodicity counterfactuals hold, but I am not even mildly confident that there is a physical probability distribution—any physical probability distribution at all—over timing schedules or deviations. I cannot regard the latter, then, as an ingredient of the former. In fact, I can make a stronger version of this argument: there is surely no physical probability distribution over the deviations in a deterministic world, because the deviations are physically impossible (hence “miracles”). Yet determinism does not deter us from asserting “likely” conditionals; thus, we must be applying a measure to deviations in these cases that does not depend in any way on physical probabilities. (This is not to say that in a fundamentally stochastic world, the measure might not take physical probabilities into account; my aim here is to show that the truth conditions for type 2 conditionals do not *require* the existence of physical probabilities, not that they ignore such probabilities when they are present.)¹⁷

15. The measure should be tweaked so that no weight is given to schedules in which two or more tosses are unrealistically close.

16. Even then, there are surely a number of reasonable choices of measure; I suggest that we evaluate type 2 conditionals by supervaluating, that is, by requiring that the consequent obtains in nearly all worlds on every such choice. The details are secondary.

17. Nevertheless, the measure over deviations *will* ignore the physical probabilities when

A final worry is that, contrary to what I have assumed so far, we are interpreting the macroperiodicity conditionals as *type 1* conditionals, that is, as asserting that under such and such counterfactual circumstances, there would be a high physical probability that the initial conditions are macroperiodic. The argument against the type 1 interpretation has the same form as the argument in the previous paragraph: I am not confident that a physical probability distribution of any sort exists over coin tosses' initial conditions, thus I am not confident that there is a fact of the matter as to whether the physical probability of macroperiodicity is high either in actual or in counterfactual circumstances, yet this in no way deters me from saying that counterfactual macroperiodicity is likely. I must therefore be asserting a type 2 conditional.

The last two paragraphs between them establish that no facts about physical probabilities go into the conditionals. In basing microconstant probability on facts about robustness, then, I am not covertly relying on some preexisting variety of physical probability: microconstant probability can be built entirely of non-probabilistic facts.

Back to the definition of microconstant probability. What remains is to build the elements of actuality responsible for the frequencies' counterfactual robustness into the definition of microconstant probability itself, as follows. The event of a system S 's producing an outcome of type e has a microconstant probability equal to p if¹⁸

1. The dynamics of S is microconstant with respect to e , and has strike

they are present. Why? Whatever physical probability a deviation has is already taken account of in the quantification its "conservatism". The measure used to evaluate the "nearly all" of a type 2 counterfactual is applied only to equally close worlds, that is, to equally conservative deviations. Thus, if the measure were to take into account the probabilities of the deviations, it would be double counting. This point is not essential to my position, but it is an aid and a comfort: it suggests that precisely the same non-probabilistic measure (non-*physically* probabilistic, that is) is used for evaluating type 2 conditionals in stochastic and deterministic worlds.

18. Here I should really say: "an outcome of type e has a microconstant probability equal to p if there exists a measure over S 's initial conditions relative to which . . ." (see notes 2 and 8).

ratio p ,

2. The actual initial conditions of nearly all long series of trials on systems of the same type as S make up macroperiodically distributed sets, and
3. The macroperiodicity of the initial conditions is robust.

Ontologically speaking, then, a microconstant probability may be thought of as constituted by three kinds of facts: the dynamical facts invoked by clause (1), the frequentist facts invoked by clause (2), and the facts responsible for the robustness invoked by clause (3). (If you want the safety of semantic ascent, think of *claims about probability as made true by three kinds of facts.*)

What are these robustness-grounding facts that form the third constituent of microconstant probability? What is microconstant probability made of, besides the dynamics in clause (1) and the actual frequencies in clause (2)? In the case of the coin toss, as demonstrated above, robustness is explained by certain dynamic facts about the physiology of coin-tossing together with some statistical information about actual pre-toss states. The same is true more generally. The macroperiodicity of the initial conditions of any microconstant process will be explained by the dynamics of the process that produces the initial conditions along with statistical facts about the relevant actual “pre-trial” states (a general term for what in the case of the coin I call pre-toss states). The third component of microconstant probability, then, consists of more dynamics and more frequencies, all readily available in a deterministic world (and none offensive to moderate empiricist sensibilities, for what that is worth).

Two remarks to conclude. The first concerns the nature of the pre-toss states. What sort of information about pre-toss states plays a part in explaining the robustness of frequencies and so goes into microconstant probability? Some answers, beginning with the most important:

1. The pre-toss information dictates the gross physiology of the human tosser or tossers.

2. It describes the natural variation of the tossers within these limits: the ways that humans differ in their physiology over time and across individuals, as well as giving a statistical profile of this variation.
3. It describes, insofar as relevant, the gross environmental conditions in which tosses occur as well as the natural variation of the environments within these limits.

This information certainly does not exhaust the details of any particular pre-toss state, but what is missing is either causally irrelevant to the determination of spin speeds or stands to be changed as a consequence of the kind of “small miracle” that brings about a typical antecedent of the robust-diagnosing counterfactuals. So the above facts alone will, I think, be sufficient, in conjunction with the relevant physical dynamics, to determine that nearly all such miracles will result in the production of macroperiodically distributed sets of spin speeds. What goes into the definition of microconstant probability is just this much in the way of actual-world facts and frequencies, and no more. (There is no need to be concerned, then, about techniques for representing the “frequencies” with which continuously varying quantities take different values.) The same goes for the generalization to any kind of pre-trial state.

My second remark concerns clause (2) of the definition of microconstant probability, the clause requiring the macroperiodicity of the initial conditions of nearly all long sequences of actual trials. I have observed that the facts that explain the robustness of macroperiodicity also explain the actual macroperiodicity itself. Why not remove clause (2) from the definition of microconstant probability altogether, then, thereby substituting for the fact of actual macroperiodicity the facts that explain it? Such a move has several advantages, of which I will mention three. First, it makes the question of how to individuate long series of trials even less important than before, since the robustness-conferring facts explain the tendency to macroperiodicity in

any long series of trials, no matter what the individuation schema.¹⁹ Second, it allows a microconstant probability to exist in a world where most sets of initial conditions for long series of trials are not macroperiodic, provided that there is a tendency to macroperiodicity and so that this actual non-macroperiodicity is no more than a fluke. Third, it allows a microconstant probability to exist in a world where there are too few actual trials to compose a long series—for example, a world in which there are only 20 coin tosses ever (provided that there are sufficiently many opportunities for tossing, that is, sufficiently many pre-toss states).

Should clause (2) be declared defunct? This depends, I think, on whether you are willing to allow that a deterministic probability could lack any predictive or explanatory value. Some will want to hold on to the idea that a probability assignment must have empirical significance. (For these thinkers, the paramount task of an account of deterministic probability is to explain the success of statistical theorizing about deterministic systems.) Some will want to hold on to the intuition of “frequency tolerance”: a probability assignment is compatible with any actual sequence of outcomes, any actual statistics. (For these thinkers, the paramount task of an account of deterministic probability is to vindicate a certain metaphysical ideal.) No story about deterministic probability can have it both ways. But I am happy to let you have it either way: if you prefer empirical significance, retain clause (2); if you prefer frequency tolerance, dispense with it.²⁰

19. I assume that all legitimate schemas attend to spatiotemporal contiguity, in the loose sense earlier described.

20. If clause (2) is removed, the resulting definition of microconstant probability might almost be taken to constitute a kind of hypothetical frequentism, except that it is not hypothetical frequencies themselves that form the basis of probability, but rather the physical dynamics and distributions of actual initial conditions that explain the frequencies.

4.4 *The Initial Condition Reference Class*

Every account of probability that indulges in a little frequentism has its “problem of the reference class” (Hájek 2007). In the case of microconstant probability, the problem manifests itself in the question: what determines the sets of actual initial conditions and/or pre-trial states, the statistics of which constitute the frequentist part of a microconstant probability?

The answer ensconced in section 4.2’s definition of microconstant probability was that the sets in question should contain the initial conditions of all long series of trials on systems of the same type. (Pre-trial states had yet to enter the picture.) On this view, it would perhaps follow (depending on your theory of explanation) that the initial conditions for spins on a normal, half-red/half-black wheel of fortune would be put in a separate class from the initial conditions for spins on an “irrational” wheel of fortune $\frac{1}{\sqrt{2}}$ of which was painted red.

Because this is an apparently invidious distinction, and in the light of the changes made to the definition of microconstant probability in the previous section, I now propose a revised solution to the reference class problem. Rather than grouping initial conditions or pre-trial states by the mechanism relative to which they are “initial” (the regular wheel, the irrational wheel, the fair coin, etc.), group them by the mechanism that produces them, or more subtly, by the factors that explain their robust tendency to assort themselves macroperiodically.

When ascribing probabilities to the irrational wheel of fortune, for example, the relevant reference class should contain systems whose initial conditions tend to be macroperiodic for the same reason that the irrational wheel’s initial spin speeds tend to be macroperiodic, that is, all systems whose initial conditions are generated in a relevantly similar way to the irrational wheel’s initial spin speeds, where “relevant” means “relevant to the explanation of macroperiodicity”. Presumably, the initial spin speeds of nearly all hand-spun wheels of fortune have a tendency to macroperiodicity for the same reasons,

and so the initial spin speeds of nearly all wheels of fortune, however they are painted, are relevant to determining the probability of *red* on the irrational wheel. Call this the *robustness rule* for determining reference classes.

The robustness rule tells you not only which initial conditions go together for the purposes of determining satisfaction of clause (2), for those of you who would like to retain that clause, but also which initial conditions are subject to clause (3), that is, which initial conditions should show a tendency to macroperiodicity, and so which facts about pre-trial states contribute to the definition of microconstant probability—namely, the kinds of facts that figure in the explanation of a certain set of systems' initial conditions' tendency to macroperiodicity.

The viability of the robustness rule turns, as you can see, on the individuation of “reasons for a tendency to macroperiodicity”, and so on the individuation of explanations of macroperiodicity's robustness. This is work for a theory of explanation (Strevens 2009), but let me make two remarks here. First, I take it that much of the work is to be done by a notion of causal-explanatory relevance. The tendencies to macroperiodicity of different sets of spins on a wheel of fortune have the same explanation because the spins are causally produced by the same kind of causal mechanism, the human body, operating under similar conditions. Second, were the causal histories of these sets of spins to be traced back in time far enough, they would no doubt diverge explanatorily in various ways: some sets of spins have their genesis in a desperate urge to gamble, others in a desire of the wheel of fortune quality control staff to do a good job, and so on. I take it that the explanatory overlap in the recent history of the sets is enough to qualify them as having a tendency to macroperiodicity “for the same reason”. In other words, the robustness rule, when comparing reasons for macroperiodicity, attends only to the proximal explanation of the tendency to macroperiodicity, the final chapter of the causal story.

Observe that the robustness rule's solution to the reference class problem

gives clause (3) a rather special flavor. It is now guaranteed by the nature of the reference class itself that the initial conditions of other systems in the class have a tendency to macroperiodicity. This aspect of clause (3) is therefore automatically satisfied. But the clause nevertheless imposes a contentful requirement on initial conditions in virtue of the presuppositions of the new definition of the reference class: in order for a microconstant probability to be ascribed to a type of system *S*, it must be that the initial conditions of every type *S* system (a) have a tendency to macroperiodicity, and (b) have this tendency for the same reason. (Whether (a) is something new depends on whether you understood the requirement that there be a tendency to macroperiodicity in the reference class to mean that there be a universal, that is, exceptionless, tendency to macroperiodicity. Exceptionlessness is now required in the systems to which the probability is to be ascribed.)

It follows from all of this that when I ascribe a probability of one-half to *red* on the spin of any wheel of fortune, I do not really mean *any* spin, but only those generated in the usual way. Most notably, a spin that is generated by some device that does not tend to produce macroperiodically distributed spin speeds is not assigned a microconstant probability by this final version of my definition. The microconstant probability of *red* is one-half, but the probability of heads on a wheel operated by a trickster who can control their spins finely enough to get *red* 75% of the time is not one-half—*red* on such a toss has no microconstant probability.

Does that mean that *red* has no probability at all? Not necessarily: there might be a probability distribution—perhaps even a microconstant probability distribution—over the trickster’s initial spin speeds. This probability distribution will, as explained in section 4.1, induce a derived physical probability for *red* (presumably roughly equal to 0.75).

I can now deliver on a promise made in section 4.1: when there is a probability distribution over a microconstant process’s initial conditions, the derived probability it induces over the process’s outcomes cannot disagree in a

serious way with the outcomes' microconstant probabilities. Why not? If the initial condition probability distribution is itself macroperiodic, the derived probabilities will equal the strike ratios and so the microconstant probabilities. If it is not macroperiodic, it will create a tendency to a non-macroperiodic distribution, in which case there will be no microconstant probabilities to disagree with. What if it is not the sort of thing that creates a tendency either way? Then we will not regard any disagreement as a difficulty: since it is the microconstant probability that is pushing the frequencies around, it is the judgments of the microconstant probability distribution and not the derived distribution that should be taken seriously.²¹

5. Conclusion

This paper has focused on gambling devices, but if microconstancy is as widespread as Strevens (2003) claims, then the definition of microconstant probability will find a place in statistical physics, population genetics, and other related branches of biology.

The systems of physics and biology are different from gambling devices in one important respect: whereas the initial conditions for a gambling device are generated by another system entirely, the human operator or croupier, physical and biological systems tend to generate their own initial conditions. The initial conditions for a month in the life of an ecosystem, for example, will be the positions and states of various organisms produced by the previous month's "trial". The situation might be compared with a gambling device

21. There is a fourth possibility: the two probability distributions (microconstant and derived) might be constructed with respect to different reference classes of initial conditions. Even in such cases, there will be a tendency for the probabilities to agree, but in general, it should not be expected that probabilities assigned to the same event with respect to different reference classes will agree. When there is disagreement, you have an epistemic reference class problem, to be solved using the resources of inductive logic rather than the metaphysics of physical probability.

in which the outcome of one trial determines the initial conditions of the next—say, a wheel of fortune with initial spin speeds marked around its circumference.

The ontology of microconstant probability is simplified in systems that generate their own initial conditions. The three parts of the definition offered in section 4.3 concerned (a) the dynamics of the system to which the probability is to be attached (e.g., the dynamics of the wheel of fortune), (b) the actual initial conditions of trials on that system (e.g., initial spin speeds), and (c) the robustness of certain properties of the actual initial conditions. The robustness is itself grounded in the dynamics and actual initial conditions of the process that generates the initial conditions. In the case of the wheel of fortune, these are facts concerning the human generation of spin speeds, but for a system that generates its own initial conditions, they are facts about the dynamics and actual initial conditions of the system itself—thus, facts of the same kind as (and largely overlapping) the facts in (a) and (b). So the ontological basis of a microconstant probability attached to such a system comprises only facts about its dynamics and facts about the distribution of its actual initial conditions.

Microconstant probability distributions are Bernoulli distributions. But some probability distributions found in the sciences are not Bernoulli—they are Gaussian, or Poisson, or Markov processes. In principle, microconstant probability can account for many of the latter kinds of probabilistic facts, however, because Gaussian, Poisson, and other probability distributions can be generated by Bernoulli processes. Whether microconstancy in fact lies at the root of the explanation of the other probability distributions is something that must be determined empirically, but such explanations are well within its power.

I conclude, then, that the definition of microconstant probability has the potential to account for all the deterministic physical probability to be found in our world. That is, deterministic probability may simply be microconstant

probability—or rather microconstant probability together with those derived probabilities that owe their existence to microconstant probability distributions over their or their ancestors' setups' initial conditions. That said, other dynamical properties may have the same power as microconstancy to explain deterministically generated statistical patterns. Perhaps these other kinds of dynamics will, by way of the techniques developed in this paper for the case of microconstancy, generate their own varieties of deterministic probability to stand alongside microconstant probability.

References

- Abrams, M. (2000). Short-run mechanistic probability. Talk given at PSA conference, November 2000. Available at <http://members.logical.net/~marshall/>.
- Bennett, J. (2003). *A Philosophical Guide to Conditionals*. Oxford University Press, Oxford.
- Bigelow, J. C. (1976). Possible worlds foundations for probability. *Journal of Philosophical Logic* 5:299–320.
- Cohen, J. and C. Callender. (2009). A better best system account of lawhood. *Philosophical Studies* 145:1–34.
- Fetzer, J. (1971). Dispositional probabilities. *Boston Studies in the Philosophy of Science* 8:473–482.
- de Finetti, B. (1964). Foresight: its logical laws, its subjective sources. In H. E. Kyburg and H. Smokler (eds.), *Studies in Subjective Probability*. Wiley, New York.
- Giere, R. N. (1973). Objective single case probabilities and the foundation of statistics. In P. Suppes, L. Henkin, G. C. Moisil, and A. Joja (eds.), *Logic, Methodology and Philosophy of Science IV: Proceedings of the Fourth International Congress for Logic, Methodology and Philosophy of Science, Bucharest, 1971*. North Holland, Amsterdam.
- Hájek, A. (1996). “Mises *Redux*” — *Redux*: Fifteen arguments against finite frequentism. *Erkenntnis* 45:209–227.
- . (2007). The reference class problem is your problem too. *Synthese* 156:563–585.
- . (2009). Fifteen arguments against hypothetical frequentism. *Erkenntnis* 70:211–235.

- Hofer, C. (2007). The third way on objective probability: A sceptic's guide to objective chance. *Mind* 463:549–596.
- Hopf, E. (1934). On causality, statistics and probability. *Journal of Mathematics and Physics* 13:51–102.
- Lewis, D. (1994). Humean supervenience debugged. *Mind* 103:473–490.
- Loewer, B. (2001). Determinism and chance. *Studies in History and Philosophy of Modern Physics* 32:609–620.
- von Mises, R. (1957). *Probability, Statistics and Truth*. Macmillan, New York. Second revised English edition, prepared by Hilda Geiringer.
- von Plato, J. (1983). The method of arbitrary functions. *British Journal for the Philosophy of Science* 34:37–47.
- Reichenbach, H. (1949). *The Theory of Probability*. University of California Press, Berkeley, CA.
- Rosenthal, J. (2010). The natural-range conception of probability. In G. Ernst and A. Hüttemann (eds.), *Time, Chance, and Reduction: Philosophical Aspects of Statistical Mechanics*. Cambridge University Press, Cambridge.
- Salmon, W. C. (1967). *The Foundations of Scientific Inference*. University of Pittsburgh Press, Pittsburgh, PA.
- Schaffer, J. (2007). Deterministic chance? *British Journal for the Philosophy of Science* 58:113–140.
- Sklar, L. (1993). *Physics and Chance*. Cambridge University Press, Cambridge.
- Sober, E. (2010). Evolutionary theory and the reality of macro probabilities. In E. Eells and J. Fetzer (eds.), *The Place of Probability in Science: In Honor of Ellery Eells*, volume 284 of *Boston Studies in the Philosophy of Science*. Springer-Verlag, Heidelberg.

Strevens, M. (2003). *Bigger than Chaos: Understanding Complexity through Probability*. Harvard University Press, Cambridge, MA.

———. (2006). Probability and chance. In D. M. Borchert (ed.), *Encyclopedia of Philosophy*, second edition. Macmillan Reference USA, Detroit.

———. (2009). *Depth: An Account of Scientific Explanation*. Harvard University Press, Cambridge, MA.