Laws About Frequencies?

John T. Roberts
johnroberts@unc.edu

For Cologne Workshop *Laws and Chances*, March 2012

Science posits a lot of probabilistic laws, including probabilistic laws of evolutionary biology, probabilistic laws of thermodynamics, and probabilistic laws of quantum mechanics, among others. The received view (among philosophers) about the probabilistic laws found in science is that they are laws about chances. "Chance" is of course a term of art; it refers to single-case objective probabilities that obey the Principal Principle (or something in the neighborhood of the Principal Principle). This view leaves us with two sets of philosophical problems concerning probabilistic laws: Problems about laws – the metaphysics of lawhood and the epistemology of laws – and problems about chances – including the metaphysics of chance and its epistemology. There's some hope that we might be able to elegantly solve both sets of problems at once – for example the way that David Lewis aims to, in his best system analysis of laws and chances together.[1] But there are two sets of problems here, and there is no guarantee that what solves one will solve the other.

Lately I've been interested in trying out an alternative to the received view: Maybe the probabilistic laws found in science are not laws about chances; maybe they are not laws about single-case probabilities at all. Maybe they are laws about frequencies instead. I mean, laws about actual frequencies, not hypothetical ones. One nice thing about this thought is that, if it's right, then the set of outstanding philosophical problems about probabilistic laws is only half as big as it appears to be. We still have the metaphysical and epistemological problems about laws, but we can get rid of the problems about chances. And unlike chances, frequencies don't bring any special metaphysical or epistemological problems of their own. So it would make life a lot easier if probabilistic laws were laws about frequencies instead of laws about chances.

---

[1] Lewis 1994.

Surely this is an idle daydream, right? The view that the objective probabilities found in science are just actual frequencies has been around for a long time, and the list of damning objections to it is even longer than the list of damning objections to the naïve regularity view of laws.

In this presentation my aim is to convince you that this idea is far less hopeless than it appears at first. I can only go so far in the available time. But I hope that I'll be able to go far enough to convince at least some of you that it's worthwhile giving this idea a run for its money.

The first thing to say is that the idea I am proposing is *not* the idea that chances are frequencies. Chances are single-case probabilities; my proposal is that there are no single-case probabilities – or rather, that if there are any, they are not what the probabilistic laws found in science are about. Instead, the probabilistic laws found in science are laws about frequencies. What's more, my proposal is not that "objective probabilities are frequencies" – the proposal actually has no use for the term "objective probabilities"; it is not a proposal about probabilities, but about probabilistic laws. So for example it does not say that what it is for the probability of getting heads on a coin-toss to be ½ is for the frequency of heads among coin-tosses to be ½.[2] *That* view is subject to a very obvious objection: It might turn out just by a complete accident that the frequency of heads is ½. Like for example if there are only ever two coin-tosses in the history of the universe, and one of them results in heads and the other in tails, but for no particular reason – that's just what happened "by chance" as non-philosophers might say. In that case, the frequency of heads would be ½, but it would be absurd to say that the *objective probability* of heads was ½ in that case. The proposal I'm

---

[2] Throughout this paper, I'm going to use coin-tosses as my main example of a type of event governed by probabilistic laws. This is only for the sake of keeping things simple. If you like, you can always substitute a more serious example without affecting the argument. For the sake of argument, I will often assume that there is a probabilistic law assigning a probability-value of ½ to heads; I intentionally refrain from making any assumptions about whether coin-tosses are genuinely indeterministic phenomena or not, since I want to allow for the possibility of probabilistic laws in a deterministic world. (More on this below.)

interested in doesn't have this problem. It doesn't say that whenever a frequency has a certain value, there is an objective probability that has that value. It just says that when there is a probabilistic law of nature, that law is rightly understood as a law about frequencies. It might be that it's a law that coin-tosses result in heads ½ of the time, in which case there's a probabilistic law about coin-tosses, but on the other hand it might just be an accidental truth that coin-tosses result in heads ½ of the time, in which case there is no probabilistic law about coin-tosses, and the frequency of heads is just an accident.

For this reason alone, the proposal I'm talking about is not just another try at the actual-frequentist interpretation of probability. It doesn't identify probabilities with frequencies; rather, it interprets those laws of nature that take the surface form of probabilistic laws as laws about frequencies, and it lets the very idea of an objective probability – as a self-standing entity that we could think about independently of thinking about the laws that govern it – fall by the wayside. So this is specifically a proposal about probabilistic laws; it is not equivalent to any interpretation of probability that doesn't mention laws. To mark this fact, I'm going to call the proposal "Nomic Frequentism."

Nomic Frequentism takes for granted the notion of a law of nature. It doesn't presuppose any particular theory of lawhood though. I think it's actually compatible with all the leading accounts of lawhood. For example, if the best-system analysis of laws is right – I mean Lewis's original best-system analysis of laws alone, not the extended best-system analysis of laws and chances together[3] – it might turn out that the best system – the one with the best combination of simplicity and strength – includes some propositions about frequencies. I think that the other leading accounts of laws could also accommodate laws about frequencies, though I don't think I'm going to have time to argue for that today.[4] So I don't think there's anything particularly Humean about Nomic Frequentism. It might be more in the spirit of Humeanism than non-Humeanism, because

---

[3] See Lewis 1994.
[4] But for details see Roberts ms, section 5 (pp. 22-26).

Humeans tend to be more skeptical about chances than non-Humeans. But in fact, you could combine the proposal that probabilistic laws are laws about frequencies with any view about laws you like – Humean or non-Humean.

No doubt, a million objections to Nomic Frequentism have occurred to you before I've even really started talking about it. I think the view has many considerable virtues, but since in my experience the idea initially strikes most philosophers as dead on arrival, I think it makes more sense to proceed by addressing some of the more obvious objections against it.

Here's a list of what I think are some of the most obvious and important objections:

1. A law about frequencies – of heads among coin-tosses, say – puts arithmetical restrictions on how many coin-tosses there can be. So for example if it's a law that the frequency of heads is ½, then its nomically necessary for the number of coin-tosses to be even. That's bizarre.

2. Frequencies aren't well-defined when the reference class is infinite or empty, but there can be probabilistic laws about classes that happen to be infinite or empty.

3. If there are laws about frequencies, then there will be "spooky action at a distance" – for example, the outcomes of coin-tosses in one region will depend on how many coin-tosses there are elsewhere and on what their outcomes are.

4. If the values of actual frequencies are constrained by laws, then this gives rise to a difficulty about probabilistic independence.

5. Laws about frequencies are about frequencies in global populations, so how do we use them to make predictions about particular cases?

6. The "Big Bad Bug" (the problem of undermining futures) is more obviously a terrible problem for frequentist views of chance than it is for other views of chance. This makes it seem as if it might be a terrible problem for Nomic Frequentism.

7. We have a strong intuition that the probabilities found in science exhibit "frequency tolerance" – they could still be what they actually are even if the frequencies were quite different. It seems that Nomic Frequentism is in conflict with this strong intuition.

8. Frequencies don't satisfy the axiom of countable additivity. But in probability theory, countable additivity has to be assumed in order to prove many important theorems. So, Nomic Frequentism hinders the ability of science to make full use of the mathematical tools of probability theory.

This paper will be structured around my replies to these eight objections. I hope that these replies will illustrate, along the way, some of the virtues of Nomic Frequentism. I will close by discussing another of its virtues: The way it makes room for probabilistic laws that do not reduce to the fundamental dynamical laws, even if the fundamental dynamical laws are deterministic. I will also suggest that Nomic Frequentism might offer an attractive alternative to the "Past Hypothesis" (Albert 2000).

FIRST OBJECTION:

> A law about frequencies – of heads among coin-tosses, say – puts arithmetical restrictions on how many coin-tosses there can be. So for example if it's a law that the frequency of heads is ½, then its nomically necessary for the number of coin-tosses to be even. That's bizarre.

Yes, that is bizarre. But it isn't really a consequence of Nomic Frequentism.

Nomic Frequentism just says that probabilistic laws are really laws about frequencies. It doesn't say that they take any particular form. The simplest form for a law about frequencies to take is exhibited by (1)[5]:

(1) $Fr(G|F) = r$

and a law of this form would indeed place arithmetic constraints on how many Fs there can be. But there are also many other forms a law about frequencies might take. For example:

(2) $Fr(G|F)$ is in the interval $[r - 1/2N_f, r + 1/2N_f]$

where N_f is the number of Fs throughout spacetime. It is easy to verify that this law places no constraints at all on how many Fs there are. It doesn't require the frequency of Gs among the Fs to be exactly equal to r, but that it be close to r – and the more Fs there are, the closer it has to be. In the limit, as the number of Fs approaches infinity, the constraint placed on the frequency converges on the constraint that it be exactly r. This answers to our natural pre-philosophical thought about what probabilistic laws are supposed to be like: If you ask a non-philosopher what it means for it to be a law that coin-tosses have a probability of ½ of landing heads, you're likely to get an answer like this: "Well, it doesn't follow that half of all coin-tosses land heads, but it does follow that about half of them will land heads – and of course the more coin-tosses there are, the closer the overall percentage of heads must be to 50%." This is exactly what the law (2) captures. And it doesn't have the bizarre consequence that there must be an even number of coin-tosses.

SECOND OBJECTION:

The same move addresses half of the second objection, namely that frequencies are undefined when the reference class is empty. That would be a problem for a law of the form (1), but it is not a problem for a law of form (2), if we naturally

[5] "Fr(G|F)" should be read as "the frequency of G among the Fs."

interpret (2) as imposing no constraints whatsoever in the limiting case as N_f approaches 0.

The other half of the second objection deals with the case where the reference class is infinite. In this case, the natural move is to say that a law about frequencies can govern a limiting frequency rather than a finite frequency. You might think this leads us into difficulties, because, for example, Richard von Mises tried to define probability in this way and ran into the problem of defining a random infinite sequence, which is a non-trivial problem.[6] But in fact, we don't run into von Mises' difficulties at all, because our aims are different than his: He was trying to give an objective interpretation of probability, whereas we're giving an interpretation of probabilistic laws.

Without getting too deeply into the details here, let me just point out that we can deal adequately with the infinite case if we recognize the possibility of one more form a law about frequencies might take.

First, I need to define a couple of terms: A "growing ball sequence" as an infinite sequence of four-dimensional spacetime balls, all centered on the same point, each one bigger than the preceding one, such that every point in spacetime is in one of the balls. A "growing ball frequency sequence for <F, G>" is a sequence f1, f2, f3 etc. where each fi is the frequency with which Fs are G within the i^{th} ball in some open ball sequence. Then a law might take this form:

(3) The limit of every growing ball frequency sequence for <F, G> is r.

In the case where there are only finitely many Fs throughout spacetime, this reduces to law (1): Fr(G|F) = r. But in the case where there are infinitely many Fs, what (3) says is perfectly well-defined, and it seems to me that it answers nicely to our natural conception of what it would be for an infinite set of Fs to conform to a statistical law. It requires that, if you look throughout a large

---

[6] See von Mises 1981; also Gillies 2000, 105-109.

enough region of spacetime, you'll find that the frequency Fr(G|F) within that region is about r – and since it quantifies over all growing ball sequences, it also requires that this goes no matter which large-enough region of spacetime you look in.

Putting (2) and (3) together, we get a form a law might take that makes sense no matter what the cardinality of the Fs is:

(4) The limit of every growing ball frequency sequence for <F, G> is in the interval $[r - 1/2N_f, r + 1/2N_f]$

(I should qualify what I just said: (4) makes sense only if the cardinality of the Fs is countable. If there are uncountably many Fs, that raises problems that I'm not going to talk about today.)

But so long as we can assume that there are at most countably infinitely many Fs, a law of the form (4) gives us just what we want: It's a plausible construal of what we might well have in mind when we suppose that there is a statistical law governing the Fs, and it doesn't place any inappropriate requirements on how many Fs there are.

I'm going to introduce a new bit of terminology: (4) can be abbreviated like this:

(5) The *-frequency of G among the Fs, i.e. Fr*(G|F), is r.

In this statement the expression 'Fr*(G|F)' is syncategorematic – properly speaking, there is no such entity or magnitude as the *-frequency of G among the Fs. But a statement that has the superficial syntactic form of an assignment of a value to this non-existent magnitude are well-defined: they are just a terminological variant on (4). So as long as we bear this in mind, it won't do any harm to talk about *-frequencies as if they were real magnitudes.

So it seems to me as if one very plausible way of spelling out the Nomic Frequentist proposal is to propose that probabilistic laws as found across the sciences be interpreted as laws that specify the values of *-frequencies.

THIRD OBJECTION: SPOOKY ACTION AT A DISTANCE

If there are laws about frequencies – for example, if there are laws of the form of (4) or (5) – then what happens in one region of spacetime depends nomologically on what happens in other regions of spacetime in a way that seems spooky. For example, suppose that there are exactly four fair coin-tosses throughout spacetime, and it is a law of nature that the *-frequency of heads-results among fair coin-tosses is ½. It follows that the frequency of heads is somewhere in the interval

[3/8, 5/8]

so, since there are four fair tosses, the frequency must be exactly ½. Now suppose that the three of the four fair tosses occurred a long time ago in a galaxy far, far away, and that to of them landed heads and one landed tails. The fourth and final fair toss is about to be performed right here and now. The laws, together with the outcomes of the other tosses entail that the outcome of this toss will be tails. This seems bizarre. Why should the outcome of this toss be dictated in this way by the outcomes of the other tosses, which occurred very far away and long ago?

In reply to the objection, the first thing to notice is that this kind of spooky action at a distance has something important in common with the more famous kind of spooky action at a distance found in quantum mechanics: Namely, you can't exploit it to send a message. In order for me to use it to send an instant message to you, which you would pick up by tossing a coin and seeing what the result was, I would have to be able to control the outcomes of the other fair coin-tosses. Assuming, what is plausible, that whatever other laws of nature there are that enable our interventions and manipulations of the environment, they do not permit us to control the outcomes of fair coin-tosses, this means I won't be able to do anything to control what result you get when you toss your coin. So, at least there is nothing here incompatible with the ban on superluminal signaling.

The second thing to notice is that the spooky action isn't quite as spooky as it might seem at first. When I presented the case of the four coin-tosses, I suggested that the outcome of the fourth toss was dictated, nomologically, by the outcomes of the other three tosses. That's how it seems to be, but that's not actually true. The law together with the outcomes of the first three fair tosses actually entail nothing whatsoever about the outcome of the fourth toss. What does fix the outcome of the fourth toss is the law, together with the outcome of the first three tosses, together with the fact that there are no other fair coin-tosses anywhere in spacetime. This extra auxiliary premise is crucial: Without it, we cannot infer that the outcome of the fourth toss will be heads. And notice how strong the extra auxiliary is: it essentially quantifies over all the events throughout spacetime. So the action at a distance here is not really the action of three coin tosses long ago and far away on a fourth one; rather, it's the action of all the rest of spacetime except the part containing the fourth coin toss, on that coin toss. It seems to me not really very spooky at all to suppose that all the rest of the contents of spacetime together with the laws of nature might determine the outcome of a fair coin-toss.

So summing up, Nomic Frequentism does imply a slightly surprising kind of action at a distance. But it's not really all that surprising, and it is far less spooky than it might appear at first.

FOURTH OBJECTION: PROBABILISTIC INDEPENDENCE (WITH A BRIEF NOTE ABOUT THE FIFTH OBJECTION: THE PREDICTION PROBLEM)

We ordinarily think that, not only does each fair coin-toss have a chance of ½ of landing heads, but also that each fair coin-toss is independent of the outcome of every other coin-toss. What's more, we think that the outcome of each fair coin-toss is independent of a lot of other things to: Whether it is raining in Paris at the time of the toss, whether the toss occurs on a Tuesday, whether the person tossing the coin has a PhD in Anthropology, and so on. If the outcome of the toss weren't independent of all those thing, and lots more, then it just wouldn't be a fair toss.

If we think that the laws governing coin-tosses are laws about chances, then we can easily make sense of this: The theory of coin-tosses would then include the following laws:

> For any fair coin-toss, the unconditional chance of a heads outcome on that toss is ½
>
> For any fair coin-toss and any possible condition that is not causally downstream of the outcome of that toss, the conditional chance of a heads outcome on that toss, given that condition, is ½.

If a Nomic Frequentist were to proceed in the same way, she would say that the correct theory of coin-tosses would include the following laws:

> The frequency of heads-outcomes among fair coin-tosses is ½
>
> For any possible circumstances C in which a coin-toss can occur – where C is not causally downstream of the outcome of the toss – the frequency of heads among fair coin-tosses in circumstances C is equal to the frequency of heads among fair coin-tosses in general, namely ½.

In short, the conditional frequency of a heads outcome on any such condition C is equal to the unconditional frequency of heads among fair coin-tosses.

Unfortunately this won't work. For example, suppose that we are about to toss a coin under circumstances that are only ever realized once in the whole history of the universe. Let 'C' be an abbreviation for those circumstances. So maybe C = it is 20:14 on a Feburary night in Cologne Germany, and John Roberts is finishing writing a paper about Nomic Frequentism. (Maybe this will never happen again.) In that case, the conditional frequency of heads on a toss in circumstances C has to be either 1 or 0. It can't be ½. So it can't be equal to the frequency of heads among coin-tosses in general.

So if we insist on understanding the relevant notion of independence as probabilistic independence, and we insist on understanding the relevant probabilities as the ones specified by the laws of nature, then the Nomic Frequentist cannot make sense of the idea that the outcome of a fair coin-toss is independent of the surrounding circumstances. But of course, a Nomic Frequentist is going to resist such insistings. She will deny that there are any such things as *the probabilities that the probabilistic laws are about*. So she will look for another way to make sense of the way in which the outcome of a coin-toss is independent of its circumstances.

Here is what I think she should say: The correct theory of fair coin-tosses says the following:

> It is a law that Fr*(Heads | Fair coin-tosses) = ½.
>
> For any type K of fair coin-toss, where K is not picked out in terms of anything causally downstream from the outcome of the toss, there is no law of nature of the form Fr*(Heads | Fair coin-tosses of type K) = r, where r is not equal to ½.

So for example, there is no law of the form "Fair coin-tosses on rainy Tuesdays in Paris come up heads with a frequency of ¾," or of any value other than ½. For all that, it may very well be true that fair coin-tosses on rainy Tuesdays in Paris come up heads with a frequency of 3/4. It's just that if that is true, then it is an accident. That's what it means to say that the outcome of a fair coin-toss is independent of whether it is a rainy Tuesday in Paris: The laws take care to regulate the frequency of heads among fair coin-tosses overall, but they don't care particularly about the fair coin-tosses on rainy Tuesdays in Paris – the laws let those fall out however they may (in the course of their making sure that overall, ½ of coin-tosses land heads). That seems a fair way of capturing the intuitive thought that the outcome of a coin-toss is independent of its circumstances, and it is available to the Nomic Frequentist.

This means that the Nomic Frequentist has to endorse theories that don't just posit laws but also quantify over laws – it has to say that there are no laws of a certain form. (My own favored theory of lawhood[7] seems to have a bit of a problem with this – it isn't perfectly obvious how my account of lawhood can handle theories that quantify over laws. Fortunately for me, that's not the topic of this paper.) And it seems we are familiar with examples of scientific theories that say exactly this sort of thing: Newton's theory, for example, is plausibly interpreted as saying that there are not any force laws that are not central-force laws. The Darwinian theory of evolution by natural selection is plausibly interpreted as saying that there are no laws to the effect that species develop in any particular teleological way, for example in the direction of greater harmony or beauty. So there isn't anything terribly novel in this suggestion that a theory might deny the existence of any laws of a certain sort.

So it seems to me that the Nomic Frequentist can, despite appearance, treat independence in a satisfactory way.

An objection presents itself here though: Suppose that we are about to flip a coin 1,000 times, and we want to make a prediction about how many of the tosses will result in heads. Surely, the right prediction to make is that about 500 of them will. But how can Nomic Frequentism deliver this prediction? Sure, it is a law that about half of all coin-tosses result in heads – but that is about the entire, global population of coin-tosses, whereas what we are interested in is the set of 1,000 tosses we are about to perform. If we knew that each individual toss had a probability of ½ of resulting in heads, and that each toss was probabilistically independent of all the others, then a straightforward calculation from

[7] In Roberts 2009.

elementary statistics would show that we should expect that about 500 of our 1,000 tosses will land heads. But the Nomic Frequentist can't say either of these things, so it seems that she cannot make use of this calculation.

What the Nomic Frequentist can say, though, is this: Suppose that we know that it is a law that half (or approximately half) of all coin-tosses land heads. And suppose we know that there is no law to the effect that coin-tosses of any special kind to which our 1,000 tosses belong will land heads with any frequency other than ½. Then we know that our 1,000 coin tosses are a sample from a huge population of tosses, approximately ½ of which land heads. On the face of it, that gives us a good reason to expect that the frequency of heads within our very large sample will be approximately ½ -- after all, almost all large samples from a population resemble the populations from which they were drawn in terms of frequencies. Of course, we would not be justified in expecting this if we had some reason to suspect that our sample was biased in some way – in other words, if we had some reason to suspect that the coin tosses in our sample belong to some subpopulation within which the frequency of heads is different from ½. But recall that we know (ex hypothesi) that there is no law to the effect that the frequency of heads within any subset of coin tosses to which our 1,000 tosses belong is anything other than ½. So, if our 1,000 coin tosses do belong to some subpopulation within which the frequency of heads is different from ½, then this fact is just a fluke. Since it is just a fluke, we cannot (presumably) be in any position to know it in advance of actually observing the results of our 1,000 tosses. So, when we are in the position of trying to make a prediction about these tosses, we are in the position of someone who has every reason to think they are dealing with a large random sample from a population within which the frequency of heads is known. A straightforward calculation taken from elementary statistics shows that we should expect about 500 of the 1,000 coin-tosses to land heads.

Cases where we use what we know about the probabilistic laws to make predictiosn about particular cases tend to be exactly like this one in the relevant respect – that is, they tend to be cases in which we are trying to make a prediction about the outcomes within a population drawn from a much larger population, where the larger population is the one that the probabilistic law is about. So there is every reason to think that the same thing goes in such cases. This shows how Nomic Frequentism accounts for the predictive role played by probabilistic laws in science.

SIXTH OBJECTION: THE BIG BAD BUG

The Big Bad Bug arises for any philosophical account of objective chance according to which *undermining futures* are possible. An undermining future is a possible future course of history F such that at the present time Ch(F) is non-zero, but in any possible world that matches our world up till now and in which F comes true, the present chances are not what they actually, presently, are. It appears that any account of chance consistent with Humean Supervenience implies that there are undermining futures. Undermining futures pose a terrible problem because, when combined with the Princiapl Principle (PP), they lead to a contradiction. Suppose that you know the way in which the chances supervene on the course of history; what should your conditional credence in F, given a specification of the actual chances and the actual history up till the present, be? By PP, it should be equal to the actual present chance of F, which is greater than zero. But given what you know, F is inconsistent with the conjunction of the history up till now and the present chances, and probability theory requires that whatever is inconsistent must have probability zero. So your credence for F should be zero, and it should be greater than zero. That is the problem.

The problem doesn't arise for the Nomic Frequentist. Consider the paradigm case of an undermining future: It's set in a world where the chance of heads on a coin-toss is ½, and it consists of a long but finite run of heads-outcomes, after the end of which there will be no more coin-tosses ever. If the chances for the different tosses are independent of one another, then the chance of this whole history would seem to be ½ raised to the power of the number of tosses in the long run, which is very tiny but not zero. According to the Nomic Frequentist, though, the probabilistic laws do not assign a chance to a long history like this.

They might say, for example, that there is a probabilistic law assigning the probability ½ to heads on a fair coin-toss, but this means that it is a law that the *-frequency of heads among fair coin-tosses is ½. They might also say that the outcome of each coin-toss is independent of all the outcomes of past coin-tosses – but this is just to say that there is no law that says that the frequency of heads on tosses following other tosses with such-and-such outcomes is anything other than ½. (Recall the discussion of independence above.) So according to the Nomic Frequentist, the correct scientific theory of coin-tosses does not assign a non-zero chance to a future consisting of an enormously long sequence of heads after which there will be no more coin-tosses ever again. Instead, it simply entails that such a future will not happen. Your credence for it, conditional on the true theory of coin-tosses, should be zero. There is no contradiction.

It might be objected that this seems like the wrong answer: It seems that we should think that it is at least possible that there will only ever be finitely many coin-tosses, and that most of them will be heads, even if we think this is enormously unlikely. And I agree: that is exactly what we should think. But that is because we should not be overly confident in the truth of any one scientific theory of coin-tosses. If Nomic Frequentism is right, then we should be completely confident that no such future will come to pass *conditional on* the hypothesis that it is a law that the probability of heads is ½. But our credence in that hypothesis should never get as high as 1. So this accounts for our lingering rational doubt about whether maybe almost all coin-tosses will come up heads after all – it is a lingering doubt that is there on account of the limitations of our epistemic situation, not because of the content of a particular scientific theory about coin-tosses that we might provisionally accept.

SEVENTH OBJECTION: FREQUENCY TOLERANCE

It seems to be a common intuition that even if there are statistical laws that specify a probability for a certain kind of outcome, the actual frequency of that outcome might diverge from that probability. Call this the intuition of "frequency tolerance."

We can distinguish a stronger from a weaker version of frequency tolerance. The weaker version says only that it is possible for the actual frequency of an outcome to differ from the probability assigned by a law – they don't have to be equal. The stronger version says that no matter what value a probabilistic law assigns to an outcome, the actual frequency of that outcome could take any value between 0 and 1: Probabilities place no logical or metaphysical or nomological limits on what value an actual frequency takes. So for example, even if the probability of getting result X from a certain type of process were 0.9999, and billions upon billions of processes of that type occur, it could still happen that none of them yield result X. It would be extremely improbable, but it would be one of those extremely improbable things that is perfectly possible.

Nomic Frequentism has no problem whatsoever with the weaker version of frequency tolerance. Laws about frequencies could take the form of *-frequency laws, or something similar, which do tolerate actual frequencies that diverge from the number occurring in the relevant law.

However, Nomic Frequentism is flatly inconsistent with the stronger version of frequency tolerance. How big a problem is this for Nomic Frequentism?

One might say here that violating the intuition of strong frequency tolerance is a liability of Nomic Frequentism, but not a decisive one, because Nomic Frequentism also has many virtues to consider, and a careful cost-benefit analysis has to be made before we draw any firm conclusions. That certainly is a wishy-washy thing to say. I think there's something much more forthright to say here: This is no problem for Nomic Frequentism at all.

Granted, many of us share the intuition that strong frequency tolerance is true. But this seems to be one of those intuitions that has everything to do with spontaneous judgments about made-up hypothetical situations, and nothing to do with anything that ever comes up in the practice of science. In scientific practice, one never accepts the proposition that a given statistical law is badly at odds with the actual, long-run, large-scale frequencies found in the world: Any indication that a putative statistical law might be that way is strong evidence against that putative law. If what we seek is a semantic analysis of the word "chance" or "probability," or a conceptual analysis of the concept *chance,* then the fact that we can so easily conceive of the probabilities and the frequencies coming apart is good evidence that we should not accept an analysis that defines probabilities as frequencies. But that is not what we are after when we do the metaphysics of probabilistic laws: There we are looking for a good interpretation of what scientific hypotheses concerning probabilistic laws come to – an interpretation that fits with all of the roles that such hypotheses play in the actual practice of science. I suggest that for an interpretation to live up to that standard, it need not respect the strong version of frequency tolerance.

The fact (supposing it is a fact) that we can clearly and distinctly conceive of there being such things as chances, which bear no logical or nomological relation to the actual frequencies at all, might show that one coherent interpretation of the statistical laws posited by scientific theories takes them to be laws about chances. But it doesn't show that this is the only possible interpretation of such laws, or that it is the best. Other possible interpretations should be judged on their ability to make sense of the role these laws play within science, rather than their ability to agree with various features of this one possible interpretation.

EIGHTH OBJECTION: COUNTABLE ADDITIVITY

I have met the objection that frequencies do not satisfy the axiom of countable additivity, which says that for any countable set of mutually exclusive propositions P1, P2, P3, ... :

$$Pr(P1 \text{ or } P2 \text{ or } P3 \text{ or } \ldots) = Pr(P1) + Pr(P2) + Pr(P3) + \ldots$$

To see why, suppose that we have a jar containing countably infinitely many marbles, each of them a different color from all the rest. Then for each of these infinitely many colors, the frequency with which that color occurs in the jar is 0 (since there is only one marble of that color, and infinitely many marbles in the jar). But the frequency with which marbles in the jar are either color 1 or color 2 or color 3 or … (including all the colors) is 1, since every single marble is one of those colors.

So frequencies do not satisfy countably additivity. This is said to be a problem for Nomic Frequentism, because many interesting theorems in probability theory depend on the axiom of countable additivity. So if Nomic Frequentism is true, that seems to threaten to limit the ability of scientists to draw on the full resources of the mathematics of probability. Anyhow, this is an objection I have heard.

But this objection is based on a very simple mistake. It is true that frequencies *need not* satisfy countable additivity. That is what the example of the jar of marbles shows. It does not follow that they *cannot* satisfy countable additivity – and indeed, they can, and they frequently do. So, there could be frequencies governed by laws that satisfied countable additivity. Moreover, there could be laws about frequencies that *entail* that the frequencies they are about satisfy countable additivity. (Why not?)

Now suppose that the laws of, say, quantum mechanics entail that the quantum probabilities do satisfy countable additivity. Then any reasonable Nomic Frequentist would interpret these laws as a set of laws about frequencies which entail, among other things, that those frequencies satisfy countable additivity. And then the full resources of probability theory could be brought to bear on quantum mechanics. On the other hand, if the laws of quantum mechanics do not, in and of themselves, entail that the quantum mechanical probabilities satisfy countable additivity, then there is something suspect about using theorems that depend on countable additivity when making applications of quantum mechanics anyway. So in no case is there any special problem about countable additivity for the Nomic Frequentist.

ON THE POSITIVE SIDE: WHAT'S GOOD ABOUT NOMIC FREQUENTISM? OR, NEED WE BE DRIVEN INTO THE ARMS OF THE MENTACULUS?

What's good about Nomic Frequentism is that it provides an interpretation of scientific hypotheses that posit statistical laws which is arguably faithful to all of the ways in which such hypotheses are used in scientific practice, without making any metaphysical assumptions whatsoever over and above whatever assumptions we have to make in order to make sense of laws of nature. So the task of explicating probabilistic laws has been reduced to the problem of making sense of laws. All the mysteries about chances disappear – in particular, the mystery of what chances could possibly be such that they obey the Principal Principle no longer need plague the metaphysics of science.

Admittedly, I have not argued here that Nomic Frequentism does make sense of all the ways in which statistical laws figure in scientific practice. That would take a much longer paper.[8] But we have seen along the way that Nomic Frequentism can make sense of applications of probabilistic laws to make predictions about actual sets of observations. Insofar as hypotheses are evaluated on the basis of checking the predictions they suggest, this suggests that Nomic Frequentism can also go some ways toward making sense of the ways in which hypotheses about probabilistic laws are evaluated in science.

Let me close by mentioning one more important virtue of Nomic Frequentism, along these lines: It can be used to make sense of probabilistic laws in a deterministic universe.

For example: Suppose our universe has a set of deterministic dynamical laws; call their conjunction D. D will be consistent with an enormous range of possible

---

[8] Roberts ms is a much longer paper which makes a start.

cosmic histories. Each of those possible histories will be nomologically possible, as far as D is concerned. But D might not exhaust all of the laws of nature that there are; there might be some additional laws as well.

For example, there might also be a law that the *-frequency with which systems of kind F have property G is r. Now, among all the possible cosmic histories associated with D, perhaps not all of them satisfy this law. In some of them, the *-frequency of G among the Fs is r, and in others it is not. So the F-G statistical law narrows down the range of nomically possible worlds to a subset of those allowed by D. Nothing in this picture threatens the determinism of the dynamical laws D. Notice also that the picture is consistent with the thesis that physics is causally/nomologically closed, in the sense that every event that occurs is or is realized by some physical event that follows from earlier physical events via the dynamical laws D.

In this way, Nomic Frequentism makes room for the possibility that there are statistical laws of the special sciences, that are not reducible to the fundamental physical laws, even in a world where the physical is causally/nomologically closed, and even in which the fundamental physical laws are deterministic. Insofar as we believe this might be a real possibility, this is a virtue of Nomic Frequentism.

The same job could be done by supposing that there is a law that narrows down the possible initial states of the universe to a range smaller than those consistent with the dynamical laws D. That is essentially the strategy that David Albert (2000) and Barry Loewer (ms) employ when they invoke the Past Hypothesis to account for the asymmetric statistical laws of thermodynamics: They make it a law of nature that the initial state of the universe is one of the ones consistent with a very low entropy at the initial moment.

But that has the disadvantage of blurring over the distinction between laws of nature and initial conditions. Some people (including me) think that distinction plays an important methodological role in science and is probably essentially bound up with our concept of a law of nature.[9]

It seems to me that adding a frequency law to the effect that the states of isolated systems are almost always entropy-increasing (or at least, non-entropy-decreasing) in one temporal direction does the same job as the Past Hypothesis – namely, that of explaining why we find temporally asymmetric thermodynamic behavior in our world even though the fundamental dynamical laws are temporally symmetric – and it does it without positing a law about initial conditions. In fact, this move doesn't even require positing that *there is* an initial moment – or even a time in the past at which the entire universe was in a state of very low entropy. For example, it could be that as we "play the film backwards," so to speak, we see isolated systems evolving into states of lower and lower entropy, so that each one eventually reaches a state of maximally low entropy and then freezes there – *but they don't all reach that maximally low-entropy state at the same time*: Systems in different parts of the universe might reach it eons apart, and there might be no time at which all of the isolated systems have already reached it. If that's how it is, then there is neither an initial state, nor a time in the past at which the entire universe was in a low-entropy state. But the statistical second law of thermodynamics still holds throughout the universe. I am not suggesting that this is a good hypothesis about the way our world is; rather the point is that we don't need the nomological necessity of the Past Hypothesis in order to establish the nomological necessity of the statistical second law, if we adopt the Nomic-Frequentist solution. This would let us say that the fact (if it is a fact) that in our universe there was an initial moment of time, and at that time the entire universe was in a very low-entropy state, is a nomologically contingent fact about the boundary conditions of our universe, though the statistical second law is indeed a law of nature. I want to suggest that this package of views might be worth considering as a plausible alternative to what Loewer (ms) calls the Mentaculus (the combination of the fundamental dynamical laws, the Past Hypothesis, a probability distribution over the possible micro-states at the initial moment).

---

[9] See also Hall ms, pp. 46-47, and p. 49 n. 21.

(This alternative package does not manage to solve what some consider to be the cardinal problem in the vicinity, namely that of explaining why we have time-asymmetric higher-level phenomena when the fundamental laws are all time-symmetric; it puts in the time asymmetry by hand, by positing a time-asymmetric law about frequencies. But of course, the Mentaculus doesn't solve that problem either: It puts in the time asymmetry by hand, by positing a time-asymmetric law about boundary conditions.)

REFERENCES:

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

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

Hall, Ned ms: "Humean Reductionism about Laws of Nature."

Lewis, David 1986: *Philosophical Papers, Volume II.* Oxford: Oxford University Press.

Lewis, David 1994: "Humean Supervenience Debugged," *Mind* 103(412): 473-490.

Loewer, Barry ms: "Two Views of Laws and Time."

Roberts, John T. 2009: *The Law-Governed Universe*, Oxford: Oxford University Press.

Roberts, John T. ms: "Laws About Frequencies." Available online at: http://philosophy.unc.edu/people/faculty/john-t.-roberts

von Mises, Richard 1981: *Probability, Statistics and Truth*, second revised English edition. New York: Dover. (Originally published in 1957 by George Allen & Unwin Ltd.)