UNDERMINED†‡

Gordon Belot
Department of Philosophy
University of Michigan

ABSTRACT. A popular strategy for understanding the probabilities that arise in physics is to interpret them via reductionist accounts of chance—indeed, it is sometimes claimed that such accounts are uniquely well-suited to make sense of the probabilities in classical statistical mechanics. Here it is argued that reductionist accounts of chance carry a steep but unappreciated cost: when applied to physical theories of the relevant type, they inevitably distort the relations of probability that they take as input.

1. Introduction

Some physical theories can be thought of as having the following form. First, the set of worlds physically possible according to the theory are singled out (typically, via equations whose solutions represent those worlds). Then, a probability measure over the space of worlds is given—call this the theory's *statistical postulate*. The paradigm example of a theory that can be put in this form is classical Boltzmannian statistical mechanics: under David Albert's influential regimentation, for example, the theory is given by specifying microphysical dynamical laws and a postulate about the thermodynamic properties of the initial condition of the universe, which jointly serve to single out a set of possible worlds, together with a further postulate that can be thought of as giving a probability measure over this set of worlds (see [Albert 2000: ch. 4; 2015: ch. 1]).

†Penultimate version. Final version to appear in the *Australasian Journal of Philosophy.*
‡Jenann Ismael, Laura Ruetsche, and Eric Winsberg aren't to blame for this paper.

The statistical postulate of a theory can seem mysterious. How should we interpret the probabilities involved? What does it mean to say that the probability of one world is *x*, while the probability of another is *y*? What would it mean to say that one such assignment was correct and another incorrect?

Some will be happy to think of the probabilities here as ideal credences. On their picture, in asserting a theory of this form, physicists first tell us to believe that the actual world belongs to a certain set of worlds, then go on to tell us how to distribute our credence over those worlds.[1]

But many find that sort of picture deeply unsatisfying. It is true that part of what you *do* when you tell me that Newton's laws of motion are true is to advise me about how manage my beliefs—but what you *say* concerns the structure of the physical world, not something about which beliefs are rational or irrational. It would be nice if we could find a way of understanding the statistical postulates that arise in physics as likewise directly concerned with the physical rather than the rational—as directly concerned not with credences but with *chances* (of, for example, particular processes of creation or annihilation of particles.[2]

In the present setting, two obstacles stand in the way. (i) In the first instance, the statistical postulate of a theory assigns probabilities to complete histories of the

---

[1] For approaches along these general lines, see, e.g., [Ismael 2013: 95 f.], [Sebens and Carroll 2014], and [Winsberg 2008: sec. 6].

[2] On traditional accounts, the content of an assertion or of a belief is a proposition (perhaps a set of possible worlds, perhaps something fancier). Moss [2015] makes a powerful case that assertions and beliefs can have as their contents sets of probability measures over spaces of propositions. On this picture, one can think of the statistical postulate of a theory as having a straightforward content, and of the assertion of that postulate as expressing the advice to bring one's credences into line with it. But even if one adopts a picture of this kind, the distinction between credences and chances remains—and many will still feel the drive to interpret statistical postulates in terms of chances.

world, rather than speaking about chances within worlds. (ii) More needs to be said about what we are talking about when we talk about chances.

The first obstacle can be overcome by noticing that assigning probabilities to possible worlds can be a way of assigning chances to possible outcomes of processes within a world. Consider simple worlds at which: time consists of just ten instants; and at each instant there are only two possible states that the world can be in. The histories of such worlds can be encoded in ten-bit binary sequences, which we will write using H's and T's rather than 1's and 0's. Suppose that we begin with all 2^{10} such histories, and then decide to assign equal probability to each history. It follows that we assign probability of .5 to the set of worlds whose history begins with H and probability .5 to the set of worlds whose history begins with T; in fact, we assign probability .5 to the set of worlds at which the state at any given instant is H (and likewise for T); and when we calculate the probability of finding any given pattern of H's and T's at any particular set of instants, we get the same answer whether or not we take into account information about what states occur at any instants outside of this set.

In this sort of case, it seems natural to say that in moving from the theory that merely told us which histories were possible to one on which each is stipulated to be equiprobable, we learn something about the chances of the state being H or T at each instant: the chances are even and they are independent of any facts about the patterns of H's and T's that occur at other instants. The state at each instant is determined by a toss of a fair cosmic coin. That, surely, is quite a different theory from what we would have if we had begun with the same set of histories, but then added a probability measure that told us that the state at each time was determined by the toss of a cosmic coin whose bias in favour of heads was .9. Progress!

What remains: explaining what talk about chance means. Some will think that our spade turns at this point—they will recommend taking chances to be metaphysically brute propensities.[3] But many prefer one or another reductionist account of chance according to which there are no such things as chances, fundamentally speaking: reality can be given a complete description in non-chancy terms; a full description of this kind licenses a family of claims about the chance facts that obtain at the world in question; but just one such family—reductionists hold that the chancy supervenes on the non-chancy, with no two worlds differing about chances without also differing in their non-chancy facts.[4]

The simplest, most direct, and least plausible form of reductionism is *frequentism*, according to which the chance of a given type of event at a world is given by the long-run frequency of such events at that world. The most popular form of reductionism is some sort of *best-system* analysis: in order to determine the laws and chances at a world, one begins with a complete description of the world in non-nomic and non-chancy terms, then looks for a package of laws (probabilistic or otherwise) that achieves an optimal combination of strength (implying as much as possible of the given description), fit (making the description as likely as possible,

[3] For propensity theories of probability, see, e.g., [Mellor 2005: ch. 4]. On the role of propensities in the interpretation of quantum mechanics, see, e.g., [Dorato and Esfeld 2010] and [Suárez 2015].

[4] This supervenience thesis does not entail reductionism—one could be a realist about chance of an epiphenomenalist stripe while holding that the chancy supervenes on the non-chancy. But because there is little to recommend such a combination, in what follows I will speak of the supervenience thesis in question as being the characteristic thesis of reductionism.

where chance is relevant), and simplicity (see [Lewis 1986: postscript C; 1994: sec. 4] and [Loewer 2001; 2004]).

So it looks like accepting some sort of reductionist account of chance allows us to tie things up in a very neat bundle—at least when the statistical postulate that we start with can be thought of as encoding facts about chance outcomes within the relevant class of worlds.[5]

The pedantic may worry about one small awkwardness in the above presentation: we start with a theory that assigns probabilities to worlds; we aim to interpret these probabilities via a reductionist account of chance; but reductionist accounts of chance aim to make sense of claims about chances at individual worlds, with no mention of theories.

My goal here is to show that the pedantic would be right to worry about this: to combine a reductionist account of chance with a theory consisting of a set of worlds equipped with a probability measure serves not (just) to interpret the probabilities that appear in that theory but (also) to change them. That is a heavy price to pay in order to uphold reductionism.

The problem we will focus on is a variant of what David Lewis [1986: 124 ff.; 1994: sec. 5] called undermining. Lewis thought undermining a merely peculiar consequence of his reductionist account of chance. But he argued that it led to a disaster once one tried to forge a link between chance and rational credence—and this argument generated the complex literature on Humean supervenience and the Principal Principle (see [Briggs 2009] for a superb critical overview). Here I don't

[5] Loewer [2001; 2004] argues that this latter condition obtains in an interesting range of cases that includes classical Boltzmannian statistical mechanics.

aim to add anything to that literature so much as to take some things away from it: my focus is squarely on chance in isolation from credence; my goal is to show just how peculiar a reductionist account of chance really is as an account of the probabilities arising in physics.

One remark about dialectic before beginning. At their most careful, advocates of best-system accounts of lawhood and chance cheerfully acknowledge that the analyses they provide have unattractive consequences—but claim that the package they offer is, all things considered, the best on the market (see [Loewer 2004; 2012]). At the same time, they claim that the package they offer is uniquely well-suited to make sense of the nature of probabilities in physical theories like classical statistical mechanics (again, see [Loewer 2004; 2012]). The aim of the present paper is to detract from the plausibility of both of these claims, by adding to the list of unattractive consequences of best-system accounts of chance the fact that they distort the relations of probability that occur in theories like classical statistical mechanics.

2. Toy Models

Consider some toy theories. In each, the state of the world can take either of two values each day; the theories differ only as to how many days there are. In one theory, T_ω, time has a beginning but no end. For each N=1, 2, 3, …, we also have a theory T_N in which time lasts for exactly N days. In any of these theories, the set of possible histories is faithfully modelled by the set of binary sequences of the appropriate length (infinite for T_ω, of length N for T_N). We will continue to write our binary sequences in terms of H and T.

A statistical postulate for such a theory is a measure that assigns probabilities to subsets of the set of sequences countenanced by the theory. We will consider only some special cases: statistical postulates that tell you to calculate probabilities of sets of binary sequences as if they were generated by flipping a coin with a bias of p in favour of heads, for some number p (strictly) between zero and one. We will denote the result of adding a measure of this kind to one of our theories by T_N^p or T_ω^p. We will focus almost exclusively on the special case of the *fair coin* measure, corresponding to p=.5.

Our question is: What content do we add in moving from T_N to $T_N^{.5}$ or from T_ω to $T_\omega^{.5}$? On a straight reading of the statistical postulate (available to non-reductionists about chance—or to anyone who just takes the postulate at face value), we entitle ourselves to say things like: the probability that the state on any given day is H is .5; and that this remains true, even if we conditionalize on information about the state on another day. More generally: the probability of getting any sort of pattern of H's and T's on any set of days is given in the obvious way by the fair-coin measure—and remains the same if we conditionalize on information about what happens on some other set of days.

So far so good. But what happens if we rely on a reductionist account of chance in explaining what all of this means? Below we will consider the finite case first, then the infinite case—in each case first seeing how things go wrong for frequentists, and then showing that the same sort of problem afflicts all forms of reductionism in virtue of their adherence to the principle that the chancy should supervene on the non-chancy.[6] Throughout, I will assume that on any reductionist reading of $T_N^{.5}$ or $T_\omega^{.5}$, the chance of the state being H on any given day is .5 (since to deny this would be to allow that the reductionist account deforms the relations of probability that it takes as input).

Before beginning, it is worth emphasizing that the problems we come across below can be expected to arise whenever a statistical postulate is read through reductionist lenses. Despite their simplicity, our toy theories are not worlds apart from more complex theories like statistical mechanics. Indeed, it is not hard to come up with problems involving boxes of gas that reduce to problems involving T_N^p or T_ω^p.[7]

3. The Finite Case: Frequentism

According to frequentism, to say that a coin is fair is to say that over the course of history heads and tails come up equally often (and more generally, that the chance of a type of outcome is p if, over the course of history, outcomes of that type occur at rate p).[8]

---

[6] So if the argument is successful, it shows in particular that best-systems approaches fare no better than frequentist approaches with respect to the problems under discussion—but it leaves untouched any account that doesn't imply the characteristic reductionist supervenience thesis.

[7] Begin with a box containing an odd number of gas molecules in thermal equilibrium. At midnight each day, divide it into two labelled compartments of equal volume. Record an H if the first compartment contains more molecules than the second compartment, otherwise record a T. The chance of H each day is .5.

[8] The considerations raised in this section are versions of traditional objections to frequentism—see, e.g., the compendium [Hájek 1997]. The main point of the present paper is to develop a theme broached (but left largely undeveloped) by Hájek—namely that more sophisticated forms of reductionism share some of the more shocking shortcomings of frequentism.

Under a straight reading, $T_N^{.5}$ says that the chance of the state being H on any given day is .5. And the same will be true on a frequentist reading of the theory.

But of course $T_N^{.5}$ says much more than that: it also says that each N-bit string corresponds to a possible history and that each of these histories is equiprobable and that the chance of the state being H on any given day remains .5 even if you conditionalize on information about the states on other days. And none of these things are true on a frequentist reading of the theory.

The problem is that frequentism on its own already tells us what the chance of H is at each world of T_N. In particular, frequentism tells us that the chance of H each day is .5 at those worlds at which H and T occur equally often over the course of history—call those the *fifty-fifty* worlds. And frequentism says that the chance of H each day is *not* .5 at worlds that are *not* fifty-fifty.

We are assuming that whatever else it says, our frequentist reading of $T_N^{.5}$ says that the chance of H on any given day is .5. And under frequentism, that rules out non-fifty-fifty sequences: the statistical postulate says that the chance of H each day is even; and under a frequentist reading, this is true at some of the worlds of T_N and false at others. So under frequentism, the move from T_N to $T_N^{.5}$ is a move from a theory that countenances worlds corresponding to all binary strings of length N to a theory that countenances just the fifty-fifty sequences. In a sense, this is good—under frequentism, imposing a statistical postulate on a theory like T_N has propositional content of the most straightforward sort (to assert the statistical postulate is to eliminate some worlds as candidates to be the actual world).

But contracting the space of worlds in this way inevitably distorts relations of probability. Consider the case N=10. Under a straight reading of $T_{10}^{.5}$, there are one thousand and twenty-four equiprobable histories; and the chance of H on Day Six is .5, no matter what the states on Days One through Five. Under the frequentist interpretation of this theory, there are only two hundred and fifty-two possible histories: more than three quarters of the histories in T_{10} are excluded as being inconsistent with the frequentist understanding of claim that the chance of H is even each day.[9] As a consequence of this contraction of the space of histories, events that are considered probabilistically independent on a straight reading of the theory come out as dependent on the frequentist reading—for example, according to a frequentist, the chance of the state being H on Day Six is zero, conditional on the state having been H on all preceding days.

These effects are generic. For any $0<p<1$ and any value of N, a frequentist reading of T_N^p involves contracting T_N to the smaller set of sequences in which H has frequency p (with the proportion of histories excluded tending towards one as N grows). This means excluding as impossible some histories that have positive probability under the straight reading of T_N^p and considering events to be probabilistically dependent that are probabilistically independent under the straight reading of T_N^p.

[9] Note that under a straight reading of this claim, each excluded history has the same probability as each history kept by the frequentist.

4. The Finite Case: Reductionism

The foregoing is just a special case of the phenomenon of undermining that drives the literature on the Principal Principle: postulate some chance laws; generate all the worlds permitted by those laws; for each world generated, run it through your favourite reductionist account of chance to see what the chances are in that world; in general you will find that your account of chance tells you that at many of the worlds generated by your initial postulate, the chance laws are inconsistent with that postulate.

Consider any account of chance that embodies the characteristic thesis of reductionism about chance, according to which at any world the chancy facts supervene on the non-chancy facts.[10] Let us call a world falling under T_N *fair* if according to this reductionist account, the chance of H at each instant is .5.[11] We continue to suppose that on a reductionist interpretation of $T_N^{.5}$, the chance of H each day is .5—so for our reductionist account, moving from T_N to $T_N^{.5}$ entails moving from the full set of worlds in T_N to the subset that count as fair according to this account.

Consider again the case N=10. There are one thousand and twenty-four histories represented by the sequences in T_{10}. It is constitutive of reductionism about chance that each of them corresponds to at most one bias p in favour of heads of that a cosmic coin might have. Now, either every ten-bit sequence is taken to

[10] It is this supervenience thesis that causes all of the trouble here and below—so dispositionalist accounts of chance and the like are not subject to the problem developed here.

[11] In general, being fifty-fifty is neither necessary nor sufficient for being fair: plausibly, for sufficiently large N best-system accounts will count some non-fifty-fifty sequences as fair (e.g., random-looking sequences in which the relative frequency of heads is some messy number very close to .5) and some fifty-fifty sequences as unfair (since, e.g., the alternating sequence HTHT…HT may be judged to describe a world at which there are no chance facts).

encode a history in which the cosmic coin is fair, or some are taken to correspond to situations in which the coin is unfair. In the former case, it becomes a necessary truth that a cosmic coin that will be tossed exactly ten times has an even chance of coming up heads on each toss—and I take it that no plausible form of reductionism can be committed to that. So on any plausible form of reductionism, some of the histories in T_{10} are unfair and must be excluded when considering $T_{10}^{.5}$.This means counting some histories as impossible which have positive probability under a straight reading of the theory.[12] And *that* means considering some events to be probabilistically dependent that are independent on a straight reading of the theory.[13]

And of course something similar holds for other values of *N* and other (non-extreme) values of *p*. Viewed through the lenses of a reductionist theory of chance, a statistical postulate for T_N is going to rule out some of the histories that that get positive probability on a straight reading of the theory—and will therefore have to distort relations of conditional probability relative to a straight reading.

(But wait! Under best-system analysis of laws without chance a phrase like 'the set of Newtonian worlds' is ambiguous between the set of worlds at which *f*=*ma*

[12] One might hope to minimize this problem by adopting a reductionist account of chance upon which only relatively few of the worlds of T_{10} count as unfair. This is not a good idea: since on reductionist accounts, each *N*-bit history falls under the statistical postulate of T_N^p for at most one *p*, one makes a reductionist account look good by bulking up the number of worlds it considers fair at the price of making it look even sillier than frequentism when it comes to its reading of T_N^p for some values of $p \neq .5$.

[13] Presumably, no plausible reductionist account will count the all-H sequence as fair—so on any such account there will be some number k=0, 1, 2,..., 9 such that if the state is H on the first k days, there is zero chance that it will be H on the next day—even though this event has probability .5 on a straight reading of $T_N^{.5}$.

and its ilk are true and the set of worlds at which they are laws.[14] One might hope that in our present setting the phase 'the set of fair coin worlds' might be similarly ambiguous—and that this ambiguity might open up a route around the difficulty identified above. But there is a big difference between the best system account of laws and the best system account of chances—a difference that renders it doubtful that such a route exists. A fully fleshed-out best-system account of (non-chancy) laws is a machine that determines which of the regularities that obtain at a world deserve the mantle of lawhood. Best-system accounts of chance cannot proceed in the same way: on reductionist accounts there are no ground-level truths about chance, so determining what the laws of chance are at a world is not a matter of simply promoting some chance-truths to chance-laws; rather, best-system accounts proceed by providing content to claims about chance at the same time as determining which hold by law (on this point, see [Lewis 1994: sec. 4] and [Loewer 2004: 1119, 1123]). So whereas in the non-chancy case there is a canonical procedure for weakening laws—the proposition that regularity R holds by law is of the form $\boxdot R$ and stripping the box off of such a proposition is a canonical way of weakening it—it is far from obvious that there is any such procedure in the chance case.)

[14] Example: a world containing a single particle moving inertially is a world at which $f=ma$ and its ilk are true, but at which they are not laws (because there are yet simpler regularities that tell you everything that there is to know about this world).

5. The Infinite Case: Frequentism

It is bad enough if the worries developed above are decisive in the finite case: much of our evidence about the behaviour of complicated systems in statistical physics and elsewhere derives ultimately from computer simulations, which can be thought of as investigating probability measures over finite spaces of possible histories.

Let us consider now the infinite case. What happens when we view $T_{\omega}^{.5}$ through frequentist lenses? Much as in the finite case, the move from T_{ω} to $T_{\omega}^{.5}$ has the effect of cutting down the set of sequences modelling worlds from the complete set of infinite binary sequences to the much smaller set of fifty-fifty sequences. But there is a disanalogy: the fair coin measure assigns probability one to the set of infinite fifty-fifty binary sequences—so in making this shift we have not flagrantly thrown away a chunk of probability as we did in the finite case.[15]

But things still go weird. Suppose that I ask you what you think the chance is that the state will be T on virtually every odd-numbered day (that is, what the chance is that the limiting relative frequency of T's on odd-numbered days is one). If you think that the histories are being generated by a fair coin, then you would be wise to say that you think that the chance is zero (whether or not you are a frequentist). But suppose that I now reveal that the state will be a H on each even-numbered day. If you are a normal person convinced that the histories are being generated by a fair coin, then you will be (very!) surprised—but will not change your odds regarding what will happen on the odd-numbered days. But if you are a frequentist convinced that the histories are being generated by a fair coin, then you

[15] There remains a measure-independent sense in which typical sequences have been excluded—see [Oxtoby 1980: 99].

will think that the chance is one that the state will be T on virtually every odd-numbered day—since that is the only pattern of H's and T's on such days consistent with the frequentist understanding of the supposition that the coin is fair.

So we again find that a frequentist reading of the supposition that chances are given by the fair coin measure rules out many histories that count as possible on a straight reading of that supposition; and, as a result, that certain events that would be independent on a straight reading fail to be independent on the frequentist reading. Here, as in the finite case, frequentism distorts the meaning of the statistical postulate.

(The above argument involves conditionalizing on an event of zero probability—something that is mathematically out of bounds on the standard formalization of probability. But there are many cases in which one knows how to perform this illicit operation.[16] So I claim that there is good reason to take arguments of this kind as having, at the very least, substantial heuristic force.[17])

6. The Infinite Case: Reductionism

Consider a reductionist account of chance. Let us again call a sequence *fair* if our account judges that relative to this sequence, the chance of an H each day is .5.

Let $\sigma=(x_1,x_2, x_3, \ldots)$ be a fifty-fifty sequence that our account considers fair. Let us divide the natural numbers into two camps: let $A:=\{k : x_k=\text{H}\}$ and let $B:=\{k : x_k=\text{T}\}$. Consider the set of S of sequences such that: (i) the reductionist account

[16] Example: if a spinner is fair, then the chance of the second spin yielding a result in the half of the wheel painted red is .5—and remains so even if we are told the precise outcome of the first spin (although that outcome, whatever it is, is an event of chance zero).

[17] I have nothing to add to the powerful case made by Hájek [2003].

under consideration considers them fair; and (ii) they agree with σ in having H's in each A-slot (in slots indexed by numbers in A).

Suppose that I tell you that a sequence falling under $T_{\omega}^{.5}$ has H's in each of its A-slots. And then I ask you what the chance is that it is fair. If you employ the reductionist account of chance under consideration to interpret $T_{\omega}^{.5}$, then you of course think that this event has chance one (since by your lights, the statistical postulate of $T_{\omega}^{.5}$ tells you that all sequences permitted by the theory are fair).

So we will end up in the same sort of problem we ran into in the preceding section—unless, that is, it turns out that if we calculate probabilities naively via the fair coin measure we also find out that a sequence with H's in each of its A slots has probability one of being fair. What would our reductionist theory of chance have to be like in order for that to be true? Take all the sequences in S and cross out the H's that appear in their A slots. This gives us a new set of sequences S^*. What sort of sequences are in S^* depends on what the reductionist account of chance we are working with looks like. Our question is: What would this account have to look like in order for S^* to have probability one according to the fair coin measure? Well—in order for any set to have positive probability according to the fair coin measure, it has to contain some fifty-fifty sequences (since the fair coin measure assigns probability zero to the set of non-fifty-fifty sequences). It follows that the only way that a reductionist account of chance can avoid counting certain events as probabilistically dependent that the fair coin measure counts as independent is if it counts some infinite sequences as fair in which the relative frequency of H's is .75. I am not ashamed to stipulate that no plausible reductionist account of chance can do that.[18]

7. Conclusion

In a recent discussion of the foundations of statistical mechanics, David Albert lets God do the talking. The setting: you have asked for a very brief but highly illuminating description of the universe—so that what you can expect to hear from God is a list of the best-system laws of our universe. God first sketches the Newtonian microphysical laws, then something like the statistical postulate of Boltzmannian statistical mechanics.

> The best I can do by way of a simple and informative description of [the initial] condition is to tell you that it was one of those which is *typical* with respect to a certain particular *probability distribution*—the Boltzmann–Gibbs distribution …. The best I can do by way of a simple and informative description of that initial condition is to tell you that it was precisely the sort of condition that you would *expect*, that it is precisely the sort of condition that you would have been rational to *bet* on, if the initial condition of the world had in fact been selected by means of a genuinely dynamically chancy procedure where the probability of this or that particular condition's being

[18] Again, it is the supervenience thesis characteristic of reductionism about chance that is causing the trouble. On non-reductionist accounts of chance, it is possible for a fair cosmic coin to generate any sequence of H's and T's. The only way that a reductionist account can achieve the same thing is by maintaining that all worlds in T_ω are fair—thus making it a necessary truth that a cosmic coin tossed an infinite number of times must be unbiased.

> selected is precisely the one given in the probability distribution of Boltzmann and Gibbs. [Albert 2015: 25 f.]

This dictum is only *something* like the usual statistical postulate of Boltzmannian statistical mechanics. For one thing, whereas the usual postulate tells you to calculate chances via the Boltzmann–Gibbs distribution, the version at hand says to do that *and* to understand chances in terms of a best-system account.[19] And like any reductionist account of chance, a best-system account will deform the relations of probability that are fed into it—for instance, some events that are probabilistically independent according to the Boltzmann–Gibbs measure will be come out as dependent on a best-system reading of the statistical postulate.

The role that the notion of typicality plays in the passage quoted above is perhaps a tipoff that Albert has more in mind than imposing a certain measure on a certain state space. Consider T_{100} equipped with the fair coin measure. On its own, this measure doesn't sort histories into the typical and the atypical. Each history is assigned the same probability, and each belongs to some sets of large measure and to some sets of small measure. More than ninety percent of one-hundred-bit sequences have between forty and sixty H's. More than ninety percent of one-hundred-bit sequences are not fifty-fifty. Is a nice, random-looking fifty-fifty sequence typical or atypical according to $T_{100}^{.5}$?

Reductionist theories of chance can be thought of as solving this 'problem': each of them can be thought of as embodying a substantive view about which

[19] For another, although Albert himself considers the fact that the initial state condition of the universe had low entropy to be a best-system law, his God never gets around to mentioning said fact. Perhaps He simply runs out of time—His prolix style is not exactly ideally suited to the task at hand.

histories are typical according to a probability measure on a state space—and as decreeing that the histories that it considers atypical are not merely improbable but impossible (relative to a given statistical postulate). This is easily seen in the case of frequentist accounts: in the case of T_{100}, for instance, adding the fair coin measure as a statistical postulate *understood in frequentist* terms means discarding more than ninety percent of the sequences falling under T_{100}—and so it is unsurprising that there are many respects in which the resulting theory departs in its judgements of probability and conditional probability from those that arise under a straight reading of the dictates of the fair coin measure.

The burden of this paper has been to argue that subtler forms of reductionism face the same sort of problem. Setting aside degenerate cases: on a reductionist account of chance, when a statistical postulate is imposed on a state space, some histories antecedently considered possible are ruled out as being inconsistent with the statistical postulate. The surviving histories will all share some feature that from the perspective of the probability measure involved in the statistical postulate is no more or less important than any other feature we could have latched on to in constructing a reductionist theory of chance. Within the newly contracted space of possibilities, relations of probability will differ from those that arise on a straight reading of the statistical postulate, available on any non-reductive reading.

Perhaps settling for this is the best we can do. But perhaps we should instead revisit non-reductive accounts of chance and accounts on which physics directly legislates credences—or follow Maudlin [2007] in exploring accounts of the role of probabilities in statistical physics that do not involve statistical postulates of the sort we have been concerned with here.

References

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

Albert, David 2015. *After Physics*, Cambridge MA: Harvard University Press.

Briggs, Rachael 2009. The Anatomy of the Big Bad Bug, *Noûs* 43/3: 428–49.

Dorato, Mauro and Michael Esfeld 2010. GRW as an Ontology of Dispositions, *Studies in History and Philosophy of Modern Physics* 41/1: 41–9.

Hájek, Alan 1997. 'Mises Redux'-Redux: Fifteen Arguments Against Finite Frequentism, *Erkenntnis* 45/2: 209–27.

Hájek, Alan 2003. What Conditional Probability Could Not Be. Synthese 137/3: 273–323.

Ismael, Jenann 2013. Naturalism on the Sydney Plan. In *Philosophical Naturalism: The Armchair or the Laboratory?*, ed. Matthew C. Haug, London: Routledge: 86–104.

Lewis, David 1986. A Subjectivist's Guide to Objective Chance, in *Philosophical Papers* Volume II, David Lewis, Oxford: Oxford University Press: 83–132.

Lewis, David 1994. Humean Supervenience Debugged, *Mind* 103/412: 473–90.

Loewer, Barry 2001. Determinism and Chance, *Studies in History and Philosophy of Modern Physics* 32/4: 609–20.

Loewer, Barry 2004. David Lewis's Humean Theory of Objective Chance, *Philosophy of Science* 71/5: 1115–25.

Loewer, Barry 2012. Two Accounts of Laws and Time, *Philosophical Studies* 160/1: 115–37.

Maudlin, Tim 2007. What Could be Objective about Probabilities?, *Studies in History and Philosophy of Modern Physics* 38/2: 275–91.

Mellor, D.H. 2005. *Probability: A Philosophical Introduction*, London: Routledge.

Moss, Sarah 2015. *Probabilistic Knowledge*, unpublished.

Oxtoby, John 1980. *Measure and Category* (second edition), Berlin: Springer–Verlag.

Sebens, Charles and Sean Carroll 2014. Self-Locating Uncertainty and the Origin of Probability in Everettian Quantum Mechanics, forthcoming in *British Journal for the Philosophy of Science*.

Suárez, Mauricio 2015. Bohmian Dispositions, *Synthese* 192/10: 3203–28.

Winsberg, Eric 2008. Laws and Chances in Statistical Mechanics, *Studies in History and Philosophy of Modern Physics* 39/4: 872–88.