Algorithmic Randomness and Probabilistic Laws

Jeffrey A. Barrett[*] and Eddy Keming Chen[†]

March 1, 2023

Abstract

We consider two ways one might use algorithmic randomness to characterize a probabilistic law. The first is a *generative chance** *law*. Such laws involve a nonstandard notion of chance. The second is a *probabilistic** *constraining law*. Such laws impose relative frequency and randomness constraints that every physically possible world must satisfy. While each notion has virtues, we argue that the latter has advantages over the former. It supports a unified governing account of non-Humean laws and provides independently motivated solutions to issues in the Humean best-system account. On both notions, we have a much tighter connection between probabilistic laws and their corresponding sets of possible worlds. Certain histories permitted by traditional probabilistic laws are ruled out as physically impossible. As a result, such laws avoid one variety of empirical underdetermination, but the approach reveals other varieties of underdetermination that are typically overlooked.

Contents

[*]Department of Logic and Philosophy of Science, University of California, Irvine, Irvine, CA 92697-5100. Email: j.barrett@uci.edu

[†]Department of Philosophy, University of California, San Diego, 9500 Gilman Dr, La Jolla, CA 92093-0119. Email: eddykemingchen@ucsd.edu

1 Introduction

Probabilistic laws, as they are usually understood, involve a variety of underdetermination. This is illustrated by a simple example.

Consider repeated tosses of a coin that produce an infinite ω-sequence of results $\langle r_1, r_2, \ldots \rangle$, where r_i is the result of the ith coin toss. Each such possible sequence of tosses gives the history of events of a possible world. Let Ω^L be the set of all such worlds that accord with a law L.

Now consider the probabilistic law L:

L**:** Each element in the ω-sequence of coin tosses $\langle r_1, r_2, \ldots \rangle$ is determined independently and with an unbiased probability of heads and tails.

One might think of L as descriptive of a fundamentally random process, something like starting with a sequence of spin-1/2 particles each in a eigenstate of z-spin, then measuring their x-spins in turn.

As probabilistic laws are typically understood, Ω^L is the set of all ω-sequences. That is, L does not rule out any world. A world compatible with L might exhibit *any* limiting relative frequency or no limiting relative frequency at all. As a result, even the full history of a world will fail to determine L in an continuous cardinality of cases.

And since Ω^L is compatible with every probabilistic law with heads and tails as possible outcomes with positive probability on each toss, even the full set of worlds compatible with L does nothing to determine L over any other probabilistic law.

This sort of underdetermination is closely related to a corresponding sort of empirical coherence.[1] A physical law is empirically coherent, in the sense we are interested in here, only if it is always in principle possible for one to have empirical support for the law if the law is in fact true.[2] If a law is empirically incoherent, then it may be impossible to learn that the law is true with even complete evidence. The law L is empirically incoherent in this sense as there are ω-sequences that might occur if L is true that would provide no empirical evidence whatsoever for accepting L. In such worlds one would never have any empirical support for accepting the correct probabilistic law even with full evidence. Indeed, since worlds compatible with L might exhibit any limiting relative frequency, there is a continuous cardinality of such worlds.

One might get a tighter fit between probabilistic laws and empirical evidence by appealing to a stronger conception of probability and a correspondingly stronger variety of probabilistic laws. Consider the law $L^\star$:

[1] See Barrett (1996), (1999) and (2020) for a presentation and discussions of empirical coherence.

[2] Throughout this paper unless specified otherwise, when we say "L is the true law" or "the law L is true," we mean not just L is true but also L is the law. This is compatible with the non-Humean perspective where laws govern and the Humean perspective where laws form the optimal description of the mosaic.

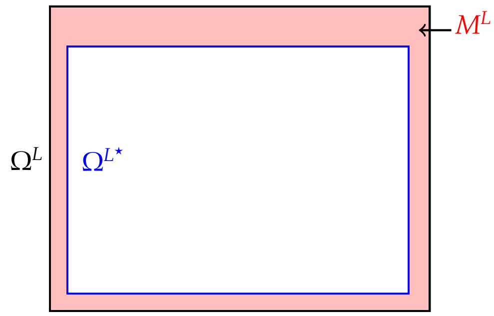

Figure 1: Ω^L, the set of worlds compatible with law L, is the set of all ω-sequences of coin toss results. $\Omega^{L^\star}$, the set of worlds compatible with law $L^\star$, is a proper subset of Ω^L. All members of $\Omega^{L^\star}$ exhibit the random pattern and relative frequencies stipulated by $L^\star$. M^L, the relative complement of $\Omega^{L^\star}$ in Ω^L, is the set of 'maverick worlds,' i.e. those that are usually regarded as compatible with L but lack the random pattern or relative frequencies.

$L^\star$**:** The ω-sequence of coin tosses $\langle r_1, r_2, \ldots \rangle$ is *random* with unbiased relative frequencies of heads and tails.

Here being *random* or not is a property of the full ω-sequence. It remains then to say what it might mean for a sequence to be random.

The notions of randomness we will consider here are algorithmic. They are defined in terms of statistical tests that determine whether a full ω-sequence exhibits any specifiable pattern. What matters at present is that each sequence will either pass or fail the test for being random.

While L is compatible with all ω-sequences of results, $L^\star$ is not. Let $\Omega^{L^\star}$ be the set of all worlds that accord with the law $L^\star$. All worlds in $\Omega^{L^\star}$ exhibit the random unbiased sequences stipulated by $L^\star$ and hence, unlike Ω^L, is a proper subset of the set of all possible ω-sequences. Specifically, $\Omega^{L^\star}$ contains no *maverick worlds*, worlds where the results exhibit a specifiable pattern or fail to exhibit the right relative frequencies or fail to exhibit any relative frequencies at all. (See Figure 1.)

If $L^\star$ is true, then *any* physically possible world fully determines $L^\star$. Further, no special probabilistic background assumptions or priors regarding what world one inhabits are required for successful inquiry.[3] A non-dogmatic inquirer in any physically possible world might determine the truth of $L^\star$ by simply conditioning on the results of coin tosses. Indeed, $L^\star$ is empirically coherent in the strong sense that, with complete evidence, one will surely learn it up to an equivalence class of computationally indistinguishable laws, something we will discuss later. A probabilistic law like $L^\star$ is much like a

[3]Both the Principle Principle and Cournot's Principle are sometimes used for this purpose. See Diaconis and Skyrms (2018, 66–7) for a brief discussion of the latter. We return to this point in §5.

deterministic law in that the law is fully determined by the evidence.

In the present paper, we consider how one might understand a ⋆-law like $L^\star$ as a *generative chance** *law* or as a *probabilistic** *constraining law*. We will then discuss the costs and benefits of such an approach. While the two notions are closely related, we will argue that thinking of $L^\star$ as a probabilistic* constraining law that governs the full sequence of coin-tosses has a number of salient virtues. They are relevant to discussions about the metaphysics of laws. As we explain in §4, the notion of a probabilistic* constraining law removes a major obstacle for developing a unified non-Humean account of governing laws, according to which laws govern by constraining physical possibilities. Such a notion also provides independently motivated solutions to the issues of the Big Bad Bug and the definition of *fit* in the Humean best-system account of laws.

2 Randomness constraints

In order to characterize a ⋆-law, one needs a test of randomness for ω-sequences. A random sequence of tosses with an unbiased coin should exhibit an even relative frequency of heads and tails in the limit. But this, of course, is not sufficient. The limiting relative frequency of an alternating sequence of heads and tails will be 1/2 for heads and tails, but this sequence is clearly not random.

There are three further conditions that one should want an unbiased random sequence to satisfy: a random sequence should be generic, patternless, and not allow for the success of a fair betting strategy.[4] These three conditions are closely related. The core idea is that an ω-sequence should count as random only if it exhibits no finitely specifiable regularity that characterizes the sequence and might consequently be used to make predictions better than chance.

Algorithmic tests are helpful in characterizing what it might mean for an infinite sequence to be patternless. As a first try, one might take an ω-sequence to be patternless, and hence random, if and only if there is no finite-length algorithm that produces the sequence.[5] If there is such an algorithm for an ω-sequence of coin-tosses, then the algorithm expresses a regularity, something that one might even think of as a deterministic law, that characterizes the sequence. But that an ω-sequence cannot be represented by a finite algorithm is again not sufficient for it to be random in the sense we are interested in here.

Consider an infinite sequence that consists of a repeated three-block pattern

[4] See Li and Paul Vitányi (2008) and Dasgupta (2011) for introductions to algorithmic complexity and randomness. See also Barrett and Huttegger (2021) for a discussion of these notions and how they relate to each other. The present section follows part of that discussion. See Eagle (2021) for an introduction to some of the philosophical issues involving randomness.

[5] One might think of an algorithm as a program in a Turing-computable language and the length of the algorithm as the length of the program. Different languages will differ in the length they assign to an abstract algorithm by no more than the length of a program that translates between the two languages.

of one thousand heads followed by one thousand tails followed by one thousand random and unbiased heads and tails. The relative frequency of heads and tails in the full sequence is unbiased. And since there are an infinite number of random blocks, such a sequence cannot be represented by a finite-length algorithm. But the sequence is clearly not random. A good Bayesian inquirer might quickly learn to bet on heads a thousand times, then bet on tails a thousand times, then bet anything at all a thousand times, then repeat the pattern. If so, she will enjoy unbounded wealth in the limit.

One way to put the problem is that there is no bound on the amount that a finite initial segment of this sequence might be compressed. One might write a very short program that takes advantage of the regularity of the blocks of heads and tails, then write a program that outputs an initial segment by alternating that short routine with a routine that just lists each random block. In this way, one might eventually shorten the algorithmic representation of finite initial segments of the sequence by more than any constant c. This observation provides the key idea behind *Martin-Löf randomness*.

An ω-sequence is Martin-Löf random if and only if there is a constant c such that all finite initial segments are c-incompressible by a prefix-free Turing machine.[6] This definition also satisfies the two other desiderata for a suitable notion of randomness. If a sequence is Martin-Löf random, then there is no fair betting strategy that generates unbounded wealth. And since measure one of infinite-length sequences are Martin-Löf random in unbiased Lebesgue measure, it meshes well with the intuition that random sequences are generic.

One might also define what it means for a sequence to be Martin-Löf random by considering the set of statistical tests that such a sequence will pass. A Martin-Löf test is a sequence $\{U_n\}_{n\in\omega}$ of uniformly Σ^0_1 classes such that $\mu(U_n) \leq 2^{-n}$ for all n, where μ is the unbiased Lebesgue measure over the sequences. Being uniformly Σ^0_1 means that there is a single constructive specification of the sequence of classes. A constructive specification is one that can be represented by a ordinary algorithm.[7] The idea is that each sequence $\{U_n\}_{n\in\omega}$ of uniformly Σ^0_1 classes corresponds to a way that a sequence might be special and thus fail an associated statistical test of randomness. A sequence passes a particular Martin-Löf test if it is not special in the specified sense.

Let 2^ω be the set of all ω-length sequences (infinite-length sequences indexed by ω). A class $C \subset 2^\omega$ is Martin-Löf null if there is a Martin-Löf test$\{U_n\}_{n\in\omega}$ such that $C \subseteq \bigcap_n U_n$. A sequence $S \in 2^\omega$ is *Martin-Löf random* if and only if $\{S\}$ is not Martin-Löf null. That is, a sequence S is Martin-Löf random if and only if it passes every Martin-Löf test. And again, a sequence has this property

[6] An initial segment is c-incompressible if and only if it is not representable by an algorithm that is c shorter than the initial segment. A prefix-free Turing machine is a universal Turing machine that is self-delimiting and hence can read its input in one direction without knowing what, if anything, comes next. Such a machine provides an even playing field. See Li and Paul Vitányi (2008).

[7] See Barrett and Huttegger (2021) for further details.

if and only if there is a constant c such that all finite initial segments are c-incompressible by a prefix-free Turing machine.

One might use the notion of Martin-Löf randomness to specify the law $L^\star$ as a constraint on the set of physically possible worlds:

$L^\star_{ML}$: The ω-sequence of coin tosses $\langle r_1, r_2, \ldots \rangle$ is *Martin-Löf random* with unbiased relative frequencies of heads and tails.

Here *all* of the worlds in $\Omega^{L^\star_{ML}}$ are random with well-defined unbiased relative frequencies. As a result, a non-dogmatic inquirer will surely infer unbiased relative frequencies in the limit. And inasmuch as all initial segments of her data will be c-incompressible, she will have as good of evidence as possible that the data are patternless and are hence randomly distributed.[8]

3 Alternative algorithmic notions

Martin-Löf randomness is not the only way that one might characterize a probabilistic coin-toss law. There are other algorithmic notions of randomness to choose from. Schnorr randomness is a closely-related notion with many of the same virtues.

A Schnorr test is a Martin-Löf test where the measures $\mu(U_n)$ are themselves uniformly computable. A class $C \subset 2^\omega$ is Schnorr null if there is a Schnorr test $\{U_n\}_{n \in \omega}$ such that $C \subseteq \bigcap_n U_n$. And a sequence $S \in 2^\omega$ is *Schnorr random* if and only if $\{S\}$ is not Schnorr null.

Schnorr randomness has similar virtues to Martin-Löf randomness. Initial segments are patternless in a strong sense, there is a natural sense in which there is no fair betting strategy, and measure one of infinite-length sequences, including all those that are Martin-Löf random, are Schnorr random. (See Figure 2.) And as with Martin-Löf randomness, one might use the notion of Schnorr randomness to specify a probabilistic constraining law:

$L^\star_S$: The ω-sequence of coin tosses $\langle r_1, r_2, \ldots \rangle$ is *Schnorr random* with unbiased relative frequencies of heads and tails.

One might also consider an associated notion of chance here as well. A chance$^\star_S$ process behaves just like an ordinary chance process except that it can never produce an infinite sequence that fails to be Schnorr random with well-defined relative frequencies.

[8] A probabilistic* law need not presume a fundamental direction of time. In order to determine an initial segment, the definition of Martin-Löf randomness seems to presuppose an initial time and a temporal direction. However, one can generalize the notion by requiring that the ordered-sequence of coin tosses be Martin-Löf random for any specified temporal direction and for any toss one regards as the "initial" toss.

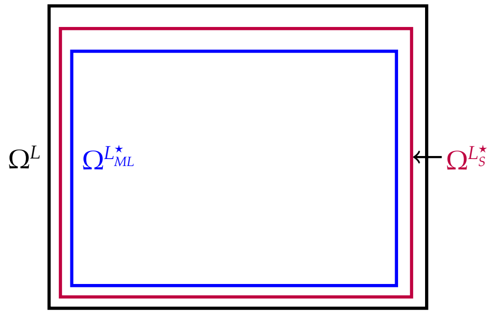

Figure 2: $L^\star_S$, which is formulated with Schnorr randomness, is compatible with more worlds than $L^\star_{ML}$, which employs Martin-Löf randomness. However, the two are empirically indistinguishable, if one is limited to Turing-strength computation.

While Martin-Löf randomness provides a particularly natural notion of randomness, Schnorr randomness also has conceptual virtues.[9] So there is a choice to make, but it is arguably a choice without empirical consequences.

Since there are sequences that are Schnorr random but not Martin-Löf random, $L^\star_{ML}$ and $L^\star_S$ are different laws. That said, they are in a strong sense empirically equivalent since there is no effective procedure that would determine whether a particular sequence is Martin-Löf random or Schnorr random but not Martin-Löf random.[10] Hence if one is limited to Turing-strength computation, one would never be able to distinguish between $L^\star_{ML}$ and $L^\star_S$ no matter what empirical evidence one had.

The upshot is that moving from a standard probabilistic law to a probabilistic constraining law eliminates one variety of empirical underdetermination, but it reveals two others. First, insofar as one expects a sequence of coin tosses governed by a traditional probabilistic law L to be such that one can detect no discernible pattern, one should expect law $L^\star$ to be empirically indistinguishable from L. And insofar as one is limited to Turing-strength computations, one will be unable to distinguish between different versions of $L^\star$ like $L^\star_{ML}$ and $L^\star_S$. Since Martin-Löf randomness has the sort of properties we want and as it is arguably the standard algorithmic notion (Dasgupta 2011), we shall understand $L^\star$ as $L^\star_{ML}$.

[9]See Downey and Griffiths (2002) for details regarding the properties of Schnorr randomness and Downey and Hirschfeldt (2010) for a description and comparison of Martin-Löf and Schnorr randomness.

[10]See Barrett and Huttegger (2021) for a proof. The notion of effective procedure here, as elsewhere in the paper when not explicitly stipulated to be otherwise, is the standard Church-Turing one.

4 Probabilities*, chances*, and laws

What kind of physical law is $L^\star$? And how does it govern the world? We suggest that one might think of $L^\star$ as a *generative chance* law* or as a *probabilistic* constraining law*. We will start with the first.

As a *generative chance* law*, $L^\star$ tells us that each toss is generated by unbiased chances*, where a chance* process behaves just like an ordinary chance process except that it can *never* produce an infinite sequence that fails to be Martin-Löf random or fails to exhibit well-defined relative frequencies.[11] As a result, a chance* process involves a subtle violation of independence. The sequence of tosses will pass every finitely specifiable test for statistical independence, but since the full sequence must exhibit the property of being Martin-Löf random with unbiased relative frequencies, a chance* process is *holistically* constrained. The constraint is not felt on any finite set of tosses, nor is it discoverable by effective means, but it does require that a relationship hold between the full sequence of tosses that is generated by the process in the limit. This interdependence between outcomes may be incompatible with the usual intuitions behind wanting a generative law. It may also be incompatible with how causal explanation works more generally.

Given this, $L^\star$ is more naturally regarded as a law that governs by constraining the entire history of the world—in this case, the full ω-sequence of outcomes. It tells us which sequences of outcomes are physically possible, namely those that satisfy the frequency constraint and the randomness constraint imposed by the law. This understanding of probabilistic laws and their governance also meshes well with Chen and Goldstein's (2022) minimal primitivism account (MinP), according to which laws are certain primitive facts that govern the world by constraining the physical possibilities of the entire spacetime and its contents.[12]

Understood this way, $L^\star$ addresses problems encountered by both non-Humean and Humean accounts of laws. We will start with the former.

On non-Humean governing accounts of laws, there is a puzzle concerning

[11] Alternatively, one might consider a similar algorithmic notion of chance but without requiring there be well-defined relative frequencies. Here we are assuming well-defined relative frequencies so that an agent might infer the law given full empirical evidence by conditioning on the results of coin tosses as she goes. The cost of this further constraint is modest since we are already requiring the full sequence to be Martin-Löf random.

[12] Three notes about the literature. (1) The present account fleshes out one of the interpretive options of probabilistic laws discussed by Chen and Goldstein (2022, §3.3.3, Option 4). (2) In one respect, the present account is similar to John T. Roberts's nomic frequentism (2009), as they both employ frequency constraints. However, Roberts does not appeal to algorithmic randomness. On nomic frequentism, non-random sequences (such as the alternating heads-tails sequence) are still physically possible. Inasmuch as any non-random sequence is regarded as evidence against the probabilistic law and in favor of a deterministic law, it would be better to exclude such sequences from physical possibilities. (3) Adlam (2022) presents an account of laws of nature as constraints that is similar to Chen and Goldstein (2022) and contains a helpful discussion of Roberts's nomic frequentism.

precisely how probabilistic laws govern. According to the standard view, probabilistic laws do not rule out any world. Instead, a probabilistic law such as L merely assigns some number between zero and one to every (measurable) subsets in the space of all ω-sequences. This raises a puzzle: what do these numbers between zero and one represent in physical reality?

Some non-Humeans appeal to gradable notions such as "propensities" (Maudlin 2007, p.20) or "probabilities of necessitation" (Armstrong 1983, p.172). Suppose a probabilistic law assigns a 0.2 probability to the next outcome being heads. On the propensities view, the chance setup has a 0.2 propensity to bring about a heads-outcome in the next toss. On the probabilities of necessitation view, the current state of affairs necessitates the state of affairs of a heads-outcome to 0.2 probability. But while one might make sense of non-gradable notions of physical possibility and impossibility, gradable notions such as propensities and degrees of necessitation are less clear. This seems undesirable.[13]

In contrast, probabilistic laws, such as $L^\star$, can be viewed as a special class of constraining laws. They constrain what is physically possible by ruling out certain sequences of outcomes, namely the maverick worlds. A sequence is physically impossible just in case it fails either the frequency constraint or the randomness constraint imposed by the law. This allows us to do away with gradable notions such as propensities or probabilities of necessitation altogether. In their place, we require only non-gradable notions of physical possibilities and impossibilities.

Consider MinP as a non-Humean example. We can now employ a single primitive relation, namely *constraining*, to understand how both probabilistic laws and non-probabilistic laws relate to the world. Both types of laws govern by constraining what is physically possible, thereby ruling out what is physically impossible. The way that $L^\star$ constrains the world is not so different from that of $F = ma$. $L^\star$ constrains the physical possibilities to be all and only the non-maverick worlds. $F = ma$ constrains the physical possibilities to be all and only the solutions of $F = ma$. In this way, $L^\star$ removes a major obstacle to a unified understanding of probabilistic and non-probabilistic laws.

Humeans may also find it useful to adopt $L^\star$ for the sort of work we have been discussing. First, it is relevant to the issue of the Big Bad Bug (Lewis 1986, pp.xiv-xvi). Lewis notices that the original version of the Principal Principle and Humean supervenience lead to a contradiction. There are certain histories of the Humean mosaic, called *undermining histories*, that are assigned, according to the Principal Principle, non-zero probability, conditionalized on some probabilistic theory T being the best system. However, they are also assigned, according to Humean supervenience, zero probability, because T would not be the best system had any of its undermining histories been actual.

Now, consider the sort of history that would count as undermining. An

[13]See Chen and Goldstein (2022, sections 2 and 3.3.3) for discussions.

undermining history either has the wrong limiting frequencies or no limiting frequencies or exhibit patterns that can be summarized by a simpler system, such as a deterministic law in the case of the alternating heads-tails sequence. Undermining histories, then, are exactly the histories of maverick worlds, as they lack the frequency or randomness patterns exemplified by typical sequences of the standard probabilistic law.

Understanding probabilistic laws as $\star$-laws rules out maverick worlds as physically impossible. If maverick worlds are physically impossible, then there are no physically possible undermining histories that can be used to derive the contradiction, and the Big Bad Bug is eliminated. Inasmuch as restricting to $\star$-laws is also motivated by considerations of underdetermination and empirical coherence, a Humean may find this solution particularly natural.

While $L^\star$ is naturally interpreted as a constraining law that is well-suited for non-Humean governing accounts such as MinP, Humeans need not interpret the constraint as something that exists over and above the mosaic. They are free to translate the present account into their preferred language by regarding $L^\star$ as a new type of Humean best-system law. They might then use it to define a new notion of Humean physical possibilities ($\Omega^{L^\star}_{Humean}$) and regard both as supervenient on the Humean mosaic.[14]

Second, Humeans who understand probabilistic laws as $\star$-laws can also avoid appealing to *fit* as a criterion in the best-system analysis of probabilistic laws, which allows them to bypass difficulties with how to characterize this notion.[15] Given an ω-sequence, there is much underdetermination among probabilistic laws such as L, and one needs something like *fit* to choose the winning best system. This is because the standard way of understanding *informativeness* as the quantity of worlds being excluded does not distinguish among probabilistic laws like L. In contrast, there is significantly less underdetermination among probabilistic laws like $L^\star$. If we consider a spectrum of different probabilistic statements like $L^\star$ that differ, say, in their specifications of the relative frequencies in the ω-sequence, then at most one of them is compatible with the ω-sequence, and thus at most one of them is an axiom in the best system of that ω-sequence. The best system analysis of a probabilistic law, such as $L^\star$, is much like that of a non-probabilistic law, such as $F = G\frac{m_1 m_2}{r^2}$. If we consider a spectrum of different versions of the Newtonian gravitational law that differ in the value of the gravitational constant G, then at most one of them is true of the mosaic, and thus at most one of them is an axiom of the best system of the mosaic. Given any Humean mosaic, one needs criteria such as simplicity and informativeness, but one does not need the statistical criterion

[14] Hoefer (2019, pp.156-158) suggests that, in his preferred solution to the Big Bad Bug, we should conditionalize on the non-occurrence of an undermining history. Thinking of the probabilistic law as $L^\star$ provides a principled reason, namely that the undermining histories are physically impossible.

[15] For helpful discussions, see Elga (2004).

of fit, to determine the best system (if there is one). Hence, the usual problems associated with fit would not arise for such Humeans.[16]

5 Discussion

We have shown how a $\star$-law may be thought of as either a generative chance* law or a probabilistic* constraining law, where the notions of chance* and probability* are subtly different from traditional chance or probability. Two differences are particularly salient.

The first concerns independence. The results of coin tosses on $L^\star$ satisfy every computable test for independence and will hence appear to be statistically independent. One might say that the results are probabilistically* independent. But inasmuch as some sequences are impossible, there is also a sense in which the results of tosses in this full ω-sequence are interdependent. To understand $L^\star$ as a generative chance* law, one would need to allow for a holistic causal structure that guarantees random sequences with unbiased relative frequencies in the limit. Depending on one's commitments regarding causal explanation, this may lead one to favor understanding $L^\star$ as a probabilistic* constraining law. If one does decide to gives up on a generative chance* law, one is, as we have just seen, left with a useful option for both proponents of governing-law accounts and Humeans.

The second difference is that chance* and probability* depend on a choice of a particular standard of algorithmic randomness. We saw this in the distinction between Martin-Löf and Schnorr randomness. But if one is limited to Turing-strength computations, there is no way to distinguish between $L^\star_{ML}$ and $L^\star_S$ on the basis of empirical evidence alone. The result is a computational sort of empirical underdetermination.[17] Since as one can only learn a $\star$-law up to an equivalence class of computationally indistinguishable laws, one might take the law $L^\star$ to be any law in this class.

The possibility of $\star$-laws reveals a further variety of underdetermination. Inasmuch as one expects the sequence of tosses one gets on a traditional probabilistic law L to be patternless, one expects L to be computationally indistinguishable from both $L^\star_{ML}$ and $L^\star_S$. Of course, these varieties of empirical underdetermination are not new, they have just gone unnoticed.

That said, $\star$-laws also help to eliminate some forms of empirical underdetermination. If $L^\star$ is true as either a generative or constraining law, then if it is among the laws that one takes seriously, then, unlike traditional probabilistic

[16]One might wonder that there is now an analogous problem of underdetermination associated with the choice between, say, $L^\star_{ML}$ and $L^\star_S$. Given this, do Humeans still need something like *fit*? It is not exactly analogous. The choice between them is not a problem for an idealized observer with the full sequence and sufficient computational power.

[17]See Barrett and Huttegger (2022) for a discussion of this point.

laws but very much like deterministic laws like $F = ma$ and $F = G\frac{m_1 m_2}{r^2}$, one will surely learn it on complete evidence in *every* physically possible world.

In contrast, if L is the true law, there will be a continuous cardinality of maverick worlds such that, if one were to inhabit any of them, one could never learn L from the results of the coin tosses. On the usual approach to thinking about laws, one needs special background assumptions to overcome this difficulty. Specifically, one needs to argue that inhabiting a maverick world of the true law is sufficiently unlikely or atypical that one has rational justification for ignoring the possibility.[18] While such assumptions may be warranted given one's other commitments, they are not required for empirical coherence if one restricts one's hypotheses to $\star$-laws.

In summary, $\star$-laws provide a way of understanding probabilistic laws as constraints on possible worlds. This helps to clarify how probabilistic laws might govern. For non-Humeans, they provide a unified way of thinking about laws as governing by constraints. And for Humeans, they provide a principled way that they might ignore undermining histories. Specifically, the new notion of probabilistic law provides a better and independently motivated way to deal with the Big Bad Bug.

6 Conclusion

We have used algorithmic randomness to characterize two types of probabilistic laws: a *generative chance** *law*, and a *probabilistic** *constraining law*. We have argued that $\star$-laws provide a novel way of understanding probabilities and chances, and help to address one variety of empirical underdetermination, but they also reveal other varieties that have been underappreciated. For all we know, our world might be characterized by a traditional probabilistic law or a $\star$-law.

In our view, the notion of a probabilistic* constraining law has advantages over that of a generative chance* law. It meshes well with the holistic character of the randomness and relative frequency constraints, directly supports a unified governing account of non-Humean laws, and provides independently motivated solutions to issues in the Humean best-system account. We suggest that both notions are worthy of study and may lead to new ideas concerning the nature of laws.

[18]Such an argument might appeal to background assumptions like the Principal Principle or Cournot's Principle. While a discussion of the status of such principles is a topic for another occasion, note that in the present case *every* world is a maverick world according to a continuous cardinality of traditional probabilistic laws. As a result, given sigma additivity, one needs to argue that one is justified in simply assigning probability zero to almost all probabilistic laws since they regard the actual world as a maverick world. This point should be uncontroversial, but it is worth emphasizing.

Acknowledgement

We are grateful for helpful discussions with Emily Adlam, Gordon Belot, Craig Callender, Eugene Chua, David Danks, Kenny Easwaran, Sheldon Goldstein, Alan Hájek, Simon Saunders, Charles Sebens, Shelly Yiran Shi, Brian Skyrms, Jason Turner, and Nino Zanghì.

Bibliography

Adlam, Emily. (2022) "Laws of Nature as Constraints," *Foundations of Physics*, **52**, 28. https://arxiv.org/abs/2109.13836

Armstrong, David M. (1983) *What is a Law of Nature?* Cambridge: Cambridge University Press.

Barrett, Jeffrey A. (2020) *The Conceptual Foundations of Quantum Mechanics*, Oxford: Oxford University Press.

Barrett, Jeffrey A. (1999) *The Quantum Mechanics of Minds and Worlds*, Oxford: Oxford University Press.

Barrett, Jeffrey A. (1996) "Empirical Adequacy and the Availability of Reliable Records in Quantum Mechanics," *Philosophy of Science* 63(1): 49–64.

Barrett, Jeffrey A. and Simon Huttegger (2021) "Quantum Randomness and Underdetermination," *Philosophy of Science* 87(3). https://doi.org/10.1086/708712

Chen, Eddy Keming and Sheldon Goldstein (2022) "Governing without a Fundamental Direction of Time: Minimal Primitivism about Laws of Nature," in Yemima Ben-Menahem (ed.), *Rethinking the Concept of Law of Nature*, Springer, pp.21-64. https://arxiv.org/abs/2109.09226

Dasgupta, Abhijit (2011) "Mathematical Foundations of Randomness," in Prasanta Bandyopadhyay and Malcolm Forster (eds.), *Philosophy of Statistics (Handbook of the Philosophy of Science: Volume 7)*, Amsterdam: Elsevier, pp. 641–710. http://dasgupab.faculty.udmercy.edu/Dasgupta-JSfinal.pdf

Diaconis, Persi and Brian Skyrms (2018) *Ten Great Ideas about Chance*, Princeton and Oxford: Princeton University Press.

Eagle, Antony, (2021) "Chance versus Randomness," The Stanford Encyclopedia of Philosophy (Spring 2021 Edition), Edward N. Zalta (ed.), https://plato.stanford.edu/archives/spr2021/entries/chance-randomness

Elga, Adam (2004) "Infinitesimal Chances and the Laws of Nature," *Australasian Journal of Philosophy*, 82(1), pp.67-76.

Hoefer, Carl (2019) *Chance in the World: A Humean Guide to Objective Chance*, New York: Oxford University Press.

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

Maudlin, Tim (2007) *The Metaphysics Within Physics*, New York: Oxford University Press.

Roberts, John T. (2009) "Laws about Frequencies," preprint, http://philsci-archive.pitt.edu/505