Quantum Randomness and Underdetermination

Jeffrey A. Barrett and Simon M. Huttegger*†

We consider the nature of quantum randomness and how one might have empirical evidence for it. We will see why, depending on one's computational resources, it may be impossible to determine whether a particular notion of randomness properly characterizes one's empirical data. Indeed, we will see why even an ideal observer under ideal epistemic conditions may never have any empirical evidence whatsoever for believing that the results of one's quantum-mechanical experiments are randomly determined. This illustrates a radical sort of empirical underdetermination faced by fundamentally stochastic theories like quantum mechanics.

1. Quantum Randomness. Randomness is a characteristic aspect of quantum phenomena. It is unclear, however, what it should mean for the results of one's quantum-mechanical measurements to be randomly distributed. It is also unclear how one might have empirical evidence for the randomness of one's measurement results. Here we will use the theory of algorithmic randomness to show how one might capture some of the standard intuitions regarding what it might mean for the results of quantum measurements to be randomly distributed. We will then see why one may never have any empirical evidence whatsoever that the results of one's quantum-mechanical experiments are in fact randomly determined even on the assumption that one's

Received February 2019; revised June 2019.

*To contact the authors, please write to: Jeffrey A. Barrett (j.barrett@uci.edu, http://faculty.sites.uci.edu/jeffreybarrett/) and Simon M. Huttegger (shuttegg@uci.edu, http://faculty.sites.uci.edu/shuttegg/), Department of Logic and Philosophy of Science, 5100 Social Science Plaza, University of California, Irvine, Irvine, CA 92697.

†This article is a direct result of a graduate seminar we taught with Sean Walsh. We would like to thank him for numerous conversations on the topic and for his insights represented by the two formal propositions. We would also like to thank Tim Maudlin, David Albert, Barry Loewer, Hannes Leitgeb, and Brian Skyrms for helpful discussions and Isaac Wilhelm, Daniel Herrmann, and the reviewers for comments on an earlier draft of this article.

Philosophy of Science, 87 (July 2020) pp. 391–408. 0031-8248/2020/8703-0001$10.00
Copyright 2020 by the Philosophy of Science Association. All rights reserved.

https://doi.org/10.1086/708712 Published online by Cambridge University Press

data are statistically uniform. This illustrates a radical sort of empirical underdetermination faced by fundamentally stochastic theories like quantum mechanics.

The argument is that once one is in the ballpark of capturing standard intuitions regarding what it might mean for the results of quantum measurements to be random, one encounters competing notions of randomness that are different but formally indistinguishable given standard computational resources. This point regarding the empirical indistinguishability of competing notions of randomness could be made more abstractly. We are considering it in the concrete context of quantum mechanics because here the truth of the physical theory depends on whether the physical world is in fact objectively random. To this end, we are interested in notions of randomness in which a physical process (as in a theory like the von Neumann–Dirac collapse formulation of quantum mechanics) or initial distribution (as in a theory like Bohmian mechanics) might be understood as being *intrinsically and objectively random*. We take this to be what is required to make sense of how physical randomness is typically understood in the context of quantum mechanics. We discuss this approach in contrast with other notions of randomness later.

In some formulations of quantum mechanics the source of quantum randomness is dynamical. This is the case for the standard von Neumann–Dirac collapse theory (von Neumann 1955) and more recent collapse theories like in Ghirardi, Rimini, and Weber (1986).[1] In other formulations quantum randomness results from the specification of special statistical boundary conditions. This is the case for some no-collapse theories. In Bohmian mechanics (Bohm 1952) quantum randomness can be thought of as resulting from the random selection of an initial particle configuration relative to the initial wave function.[2] In other no-collapse formulations quantum randomness is the result of epistemic uncertainty regarding self-location. This is the case for some many-world reconstructions of Everett's pure wave mechanics (Dirac 1957).[3] Here we will consider quantum randomness in the context of the standard von Neumann–Dirac formulation of quantum mechanics, but these considerations are also applicable to other formulations of quantum mechanics that appeal to the notion of a random process or random selection.

The standard von Neumann–Dirac collapse formulation of quantum mechanics stipulates that one's measurement results are the outcome of a random dynamical process and, hence, predicts that a sequence of measurement

1. See Albert (1992) and Barrett (2019) for discussions.
2. See Barrett (1999, 2019) for discussions.
3. See Saunders et al. (2010), Wallace (2012), and Barrett (2019) for discussions.

https://doi.org/10.1086/708712 Published online by Cambridge University Press

results will be randomly distributed.[4] We first consider the sense in which it predicts random measurement results and then consider how one might empirically test its predictions.

On the standard collapse formulation of quantum mechanics, the state of a physical system S is represented by an element $|\psi\rangle_S$ of unit length in a Hilbert space $\mathcal{H}$, and a physical observable O is represented by a Hermitian operator $\hat{O}$ on that space. The physical interpretation of a state is given by the eigenvalue-eigenstate link, which says that a system S has a determinate value for observable O if and only if it is in an eigenstate of O. That is, S has a determinate value for O if and only if $\hat{O}|\psi\rangle_S = \lambda|\psi\rangle_S$, where $\hat{O}$ is the Hermitian operator corresponding to O, $|\psi\rangle_S$ is the vector representing the state of S, and the eigenvalue λ is a real number. In this case, one would with certainty get the result λ if one measured O of S.

Given the eigenvalue-eigenstate link and the linear way that systems evolve when they are not measured, a particular observable will typically fail to have any determinate value at all for a given system before the system is measured. According to the standard theory, the system *acquires* a determinate value for the observable when it is measured. In particular, the theory predicts that when the observable is measured, the system will instantaneously and randomly jump to an eigenstate of the observable being measured with probabilities determined by its initial state. Since the final state will be an eigenstate of the measured observable, it will be one where the object system now has a determinate value for that observable. And, salient to the present discussion, that value will be *randomly determined* by the process that generated it.

In describing the dynamical laws of the standard theory, von Neumann referred to the random nonlinear evolution of the state that occurs on measurement as process 1. When no one is observing the system, it evolves in a deterministic, linear way that he called process 2. These two dynamical laws might be characterized as follows:

> Process 1: If a *measurement* is made of the system S, the state of S will *randomly collapse* to an eigenstate of the observable being measured (a state where the system has a determinate value of the observable being measured). If the initial state is given by $|\psi\rangle_S$ and $|\chi\rangle_S$ is an eigenstate of O, then the probability of S collapsing to $|\chi\rangle_S$ is equal to $|\langle\psi|\chi\rangle|^2$ (the square of the magnitude of the projection of the premeasurement state onto the eigenstate).

4. See Barrett (2019) for a discussion of the standard von Neumann–Dirac formulation of quantum mechanics and its conceptual structure.

https://doi.org/10.1086/708712 Published online by Cambridge University Press

> Process 2: If *no measurement* is made of a physical system, it will evolve in a deterministic, linear way: if the state of S is given by $|\psi(t_0)\rangle_S$ at time t_0, then its state at a time t will be given by $\hat{U}(t_0, t_1)|\psi(t_0)\rangle_S$, where $\hat{U}$ is a unitary operator that depends on the energy properties of S.

Process 2 is a deterministic dynamical law that explains quantum-mechanical interference and entanglement. In contrast, process 1 is a fundamentally stochastic dynamical law. It explains both why measurements yield determinate outcomes and why one should expect a sequence of quantum measurement results to be *randomly distributed with the standard quantum statistics*.

That the theory does not say what constitutes a *measurement* means that it is unclear precisely when each dynamical law obtains. This ambiguity is the source of the quantum measurement problem.[5] For present purposes, we will simply suppose that process 1 kicks in at some point during a measurement interaction to produce determinate measurement records that are randomly determined with the standard quantum statistics. Our concern here is not to say precisely when or why collapses occur but rather to consider what it might mean to say that one's measurement records are randomly determined with the standard quantum statistics and how one might have empirical evidence for such a claim.

Consider an infinite series of systems S_1, S_2, ... S_k, ... each in the state

$$1/\sqrt{2}(|\uparrow_x\rangle_{S_k} + |\downarrow_x\rangle_{S_k}).$$

Suppose that one measures the x-spin of each system in turn and records 0 for $\downarrow_x$ and 1 for $\uparrow_x$ as a string σ. Call this the *quantum coin-toss experiment*.[6] Here process 1 predicts that the outcome of each trial will be randomly determined with probability 1/2 of recording 0 and probability 1/2 of recording 1 on each trial. To say something that is at least very closely related, one expects with probability one that the outcomes to be statistically independent and unbiased.[7]

5. See Albert (1992) and Barrett (1999, 2019) for discussions of the quantum measurement problem and various proposed resolutions.

6. One might equivalently, according to the standard theory, alternate x-spin and z-spin measurements on a single electron and keep track of the sequence of up and down results.

7. There is an important distinction to be made between a *random sequence* and a sequence produced by a *random process* as these notions are typically used. While one would expect (with probability one) a random sequence from a random process, and a random sequence is empirical evidence that it was generated by a random process, it is possible for a random process to produce a nonrandom sequence. But if a process does produce a nonrandom sequence, that clearly counts as evidence against the process being random inasmuch as that would be extraordinarily unlikely otherwise.

https://doi.org/10.1086/708712 Published online by Cambridge University Press

While no particular sequence of 0's and 1's is ruled out in the quantum coin-toss experiment, one would have very good empirical evidence against process 1 if the ratio of 0's and 1's in σ were not approximately even in the long run. If so, the dynamics would be predicting the wrong relative frequencies. But one would also have very good empirical evidence against process 1 if the sequence of results exhibited a simple pattern like 01010101. . . . This would not count against the dynamics predicting the right relative frequencies, but the longer a simple pattern like this persists, the better one's empirical evidence that the measurement results are not statistically independent and hence not in fact determined by a random process at all.

If process 1 is descriptive of the physical world, then one should expect a sequence of measurement results that exhibits all of the properties of a random sequence. One such property in the present case is that a random sequence of measurement results should be expected to have the standard quantum relative frequencies. But, as the example of an alternating sequence of zeroes and ones illustrates, having the right relative frequencies is not sufficient for the sequence of measurement results to be randomly distributed. We expect the sequence to exhibit other statistical features as well. That said, it is not immediately clear what these should be. In addition to tracking relative frequencies, one needs an explicit test of all of the features of a random sequence, whatever these may be, in order to check the empirical predictions of process 1.

The general methodological question here concerns how one might empirically determine whether the output of a physical process is in fact random. Equipped with a test for randomness, a good Bayesian might then seek to update her degree of belief that the sequence σ was produced by a random process by conditioning on new measurement results as one gets them. But what should it mean for the results in σ, or an initial segment of σ, to be random?

As suggested earlier, our intuitions concerning what it means for a sequence of results to be randomly distributed are closely tied to our intuitions concerning what it means for those results to be statistically independent. The judgment that the sequence 01010101. . . does not appear to be random goes hand in hand with the judgment that the measurement results that constitute it do not appear to be statistically independent. In this sense, a test for statistical independence is a test for a corresponding variety of randomness and the other way around.[8] If process 1 is in fact descriptive of the physical world, then one should expect the results of the quantum coin-toss experiment

8. Events A and B are *statistically independent* if and only if $P(A) = P(A|B)$. The issue here is how one tests whether this condition is satisfied by the dynamical process that produced one's results given those results.

https://doi.org/10.1086/708712 Published online by Cambridge University Press

to be both random and statistically independent in some appropriately strong sense.

Von Neumann's physical intuition was that the sequence σ should be expected to be random and its elements statistically independent because it is determined by a dynamical process that produces *arbitrary* events.[9] Specifically, he understood process 1 to postulate a *willkürliche Veränderung*—an arbitrary or capricious change in the physical state. Because the sequence of measurement results is arbitrary, one expects it to be patternless and not special in any specifiable sense. The arbitrary results of this process were, for von Neumann, what made the standard quantum probabilities the most precise empirical predictions possible. He took the quantum mechanical state to be complete. Further, salient for the issue at hand, he took process 1 to be (1) a physical process and (2) dynamically complete—there is simply nothing more to say about what determines the result of a quantum measurement.[10]

As a consequence of von Neumann's commitment to the outcomes being arbitrary, the sequence σ should be *typical*. That is, it should be a sequence that one can think of as having been arbitrarily selected from a subset of measure one of all possible sequences in Lebesgue measure, just as in the case of the random outcomes of tosses of a theoretical fair coin. This ties directly to statistical independence. If the measurement results are independent, then one should also expect the sequence to be typical in Lebesgue measure.

Finally, concerning von Neumann's commitment to state completeness and the completeness of the dynamics, since the standard probabilistic predictions of quantum mechanics are the most precise predictions possible, one should expect there to be no fair betting strategy that would allow one to do better in predicting the sequence σ than simply predicting each result with the standard quantum probabilities—here each with probability 1/2.[11]

2. Algorithmic Randomness. In order to test the empirical predictions of the standard theory, then, we want to test σ for being patternless in a way that satisfies our statistically independent, unpredictable, no-betting-strategy intuitions. In short, we want to ensure that the sequence exhibits no specifiable regularity.

9. By *expectation* we mean expectation with probability one.

10. See Wigner (1970, 1005–6 and n. 1), Earman (1986, 227–28), and Barrett (2019) for more on von Neumann's understanding of the essential nature of quantum unpredictability. See Earman (1986, 199–234) for a discussion of the distinction between quantum indeterminism and quantum unpredictability.

11. A fair betting strategy is a plan to bet for or against outcomes at each stage so that one does not expect to win or loose at the next stage. If one were to adopt an unfair betting strategy, one could of course expect to win arbitrarily large amounts of money.

https://doi.org/10.1086/708712 Published online by Cambridge University Press

In this spirit, one might take a sequence to be patternless, and hence random, when there is no algorithm significantly shorter than the sequence that produces it. While this is a step in the right direction, it is not quite what we want. An immediate problem is that an infinite sequence might be incompressible but still contain long, recurring subsequences that exhibit regular patterns. In that case, the sequence as a whole does not satisfy our intuitions regarding what it is to be random.

Consider an infinite sequence that consists of one thousand 1's followed by one thousand 0's followed by one thousand random 0's and 1's, then repeats this three-block pattern forever. Because of the random blocks, such an infinite sequence may be incompressible in the sense of not being representable by a finite-length algorithm, but the full sequence is clearly not random. This is reflected by the fact that there is a simple betting strategy that would lead to unbounded wealth in the long run (e.g., predict 1 a thousand times, then 0 a thousand times, then whatever one wants a thousand times and repeat). The upshot is that this very simple notion of algorithmic randomness is too weak to support the intuition that there should be no pattern or betting strategy that allows one to predict better than chance. But we are on the right track.

There are a number of more subtle notions of algorithmic randomness that do support our patternless, statistically independent, unpredictable, no-betting-strategy intuitions. We will consider two here: *Martin-Löf randomness* and *Schnorr randomness*. Both of these satisfies the basic intuition that a random sequence should be patternless in a way that makes it effectively unpredictable and, in a strong sense, does not allow for a successful betting strategy.

Martin-Löf and Schnorr randomness fit into a hierarchy of algorithmic ways of understanding what it might mean for a sequence to be random. The core notions of algorithmic randomness from less to more restrictive are Kurtz (weak) random, Schnorr random (SR), computably random, Martin-Löf random (MLR), and 2-random. We are concerned here with Schnorr randomness and Martin-Löf randomness.[12]

A notion of randomness can be given in terms of a set of tests that a random sequence will pass. A Martin-Löf test is a sequence $\{U_n\}_{n\in\omega}$ of uniformly Σ^0_1 sets of sequences 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 sets. A constructive

12. See Li and Vitányi (2008) for an introduction to algorithmic complexity and randomness. See Downey and Hirschfeldt (2010) for a description and comparison of Martin-Löf and Schnorr randomness and Downey and Griffiths (2002) for further details regarding the properties of Schnorr randomness.

https://doi.org/10.1086/708712 Published online by Cambridge University Press

specification can be represented by a ordinary algorithm.[13] Let 2^{ω} be the set of all ω-length sequences (infinite-length sequences indexed by ω). A sequence $S \in 2^{\omega}$ passes the Martin-Löf test associated with the sequence $\{U_n\}_{n\in\omega}$ if and only if S is not in the measure-zero null set $\cap_n U_n$. Passing a Martin-Löf test can be thought of as passing a particular effective statistical test for randomness.[14] A sequence is MLR if and only if it passes every Martin-Löf test. Since there are only a countable number of Martin-Löf tests, the union of all the associated null sets is also a set of measure zero. So the set of MLR sequences has Lebesgue measure one.

The idea here is that each sequence $\{U_n\}_{n\in\omega}$ of uniformly Σ^0_1 classes corresponds to an effectively specifiable way that a sequence might be special and thus to an associated statistical test of randomness. A sequence passes the test if it is not special in the specified sense. A sequence is MLR then if (1) it is not special in any way that can be effectively specified, and hence (2) it passes every effective statistical test for being random. This is arguably precisely what one should want for a sequence to be considered random.

Martin-Löf randomness also supports the intuition that a random sequence should be patternless in the sense of being both incompressible and unpredictable. An infinite sequence is MLR if and only if there is a constant c such that all finite initial segments are c-incompressible (not representable by an algorithm that is c shorter than the initial segment) by a prefix-free machine (a universal Turing machine that is self-delimiting and hence can read its input in one direction without knowing what, if anything, comes next). And a sequence is MLR if and only if no constructive martingale succeeds on it (if there is no constructive betting strategy that generates unbounded wealth).[15] Since measure one of infinite-length sequences are MLR in the unbiased Lebesgue measure over the set of possible sequences, it also supports the intuition that random sequences are not special in a measure-theoretic sense.

The notion of a sequence being SR is closely related. A Schnorr test is a Martin-Löf test where the measures $\mu(U_n)$ on the sequence of uniformly Σ^0_1 sets are uniformly computable (there is a single algorithm that computes each of these measures). A sequence $S \in 2^{\omega}$ passes the Schnorr test associated with the sequence $\{U_n\}_{n\in\omega}$ if and only if S is not in the measure-zero null set $\cap_n U_n$. A sequence is SR if and only if it passes every Schnorr test. Because the measures on the test classes $\mu(U_n)$ are uniformly computable,

13. The Σ^0_1 sets are semicomputable open sets in the following sense. Every Σ^0_1 is the union of a countable set of cylinder sets, the clopens of Cantor space. By taking an increasing sequence of finite unions of clopens, we can approximate each Σ^0_1 set by computable objects from below.

14. See Downey and Hirschfeldt (2010, 229–31) for an extended discussion.

15. *Constructive* again means computably approximable from below.

https://doi.org/10.1086/708712 Published online by Cambridge University Press

the statistical tests here might be thought of as being more concretely specifiable than in the case of Martin-Löf randomness. Indeed, one can suppose that the measures of the test classes are given by $\mu(U_n) = 2^{-n}$ without loss of generality.

The notion of Schnorr randomness, like that of Martin-Löf randomness, captures the intuition that a random sequence should be patternless in the sense of being both incompressible and unpredictable in a strong sense. An infinite sequence is SR if and only if there is a constant c such that all finite initial segments are c-incompressible by a computable measure machine (a prefix-free Turing machine with a domain of computable measure).[16] If a sequence is SR, then no computable martingale h-succeeds on it (there is no computable betting strategy that generates wealth over time that is bounded from below by an unbounded, nondecreasing function h).[17] And, like MLR sequences, SR sequences are not special—measure one of infinite-length sequences are SR.

Important for what follows, MLR infinite-length sequences are a proper subset of SR sequences. Since MLR sequences and SR sequences are both measure-one sets, their intersection is also measure one. And the set of sequences that are SR but not MLR is measure zero in the unbiased Lebesgue measure over the set of infinite-length sequences. Sequences that are SR but not MLR are measure theoretically very special.

3. The Effective Indeterminacy of Randomness and Independence. Given how they support the relevant intuitions, both MLR and SR provide plausible standards for quantum randomness. Indeed, inasmuch as random sequences are not special, one would expect (with probability one) the sequence σ of quantum-mechanical results produced by process 1 to be both MLR and SR. But here one encounters a number of epistemic problems.

Consider the following proposition concerning whether one can know whether a sequence is MLR or SR.[18]

> **Proposition 1.** Suppose that $\mathcal{C} \subseteq 2^{\omega}$ is a nonempty class such that either (i) $\mathcal{C}$ contains no computable members or (ii) $\mathcal{C} \neq 2^{\omega}$ and $\mathcal{C}$ is a tail set; that is, if X is in $\mathcal{C}$ and Y differs from X by at most finitely many bits, then Y is in $\mathcal{C}$. Then there is no algorithm e such that for all $X \in 2^{\omega}$ one has $\varphi_e^X(0) = 1$ if

16. See Downey and Hirschfeldt (2010, 277) for further details on such machines.

17. See Downey and Hirschfeldt (2010, 271) for further details regarding the martingale properties of SR sequences.

18. See Soare (2016) for an explanation of the notation here. The proof of this proposition follows closely from the definitions of the relevant notions. See, e.g., Soare (2016, 190) and Shen, Uspensky, and Vereshchagin (2017, 81). See also Downey and Hirschfeldt's (2010, 16–18) discussion of the use principle.

https://doi.org/10.1086/708712 Published online by Cambridge University Press

and only if $X \in \mathcal{C}$, where φ_e^X denotes the eth computable function with oracle X.

Proof. Suppose not, with witness e. Since $\mathcal{C}$ is a nonempty class, choose X in $\mathcal{C}$. Then $\varphi_e^X(0) = 1$. Then there is s such that $\varphi_{e,s}^X(0) = 1$; that is, the computation converges in $<s$ steps looking at $<s$ bits of the oracle tape.
Let $\sigma = X \upharpoonright s$. Suppose that i is satisfied. Consider $Y = \sigma^\frown \bar{0}$ (i.e., σ followed by all zeros). This is computable, and we also have that $\varphi_{e,s}^Y(0) = 1$ and hence $Y \in \mathcal{C}$, contradicting our hypothesis that $\mathcal{C}$ contains no computable members.
Suppose that ii is satisfied. Then for any $Y \in [\sigma]$, that is, any Y that begins with σ, we have $Y \in \mathcal{C}$. But every element of 2^ω differs by an element of $[\sigma]$ by only finitely much, and since $\mathcal{C}$ is a tail set, we then have $\mathcal{C} = 2^\omega$.

Notions of algorithmic randomness typically satisfy both conditions i and ii. In particular, both Martin-Löf and Schnorr randomness satisfy these two conditions. The upshot is that there is no effective procedure to tell whether a sequence σ is MLR or is SR. This means that if one is restricted to Turing-strength computations, one can never know whether one's empirical evidence is in fact random in either of these two senses.[19] But the epistemic situation is significantly worse than this might suggest.

The following proposition is concerned with the question of whether one can distinguish between sequences that are SR but not MLR and sequences that are MLR.

Proposition 2. There is no algorithm e such that for all $X \in 2^\omega$, if one has that if X is SR, then $\varphi_e^X(0) = 1$ if and only if X is MLR.

Proof. Choose MLR X. Then as above, $\varphi_{e,s}^X(0) = 1$ for some s, and again set $\sigma = X \upharpoonright s$. Choose Y that is SR but not MLR, and let $Z = \sigma^\frown Y$. Then since the SRs and the MLRs are tail sets, one has that Z is SR but not MLR. But we also have that $\varphi_e^Z(0) = 1$ since $Z \in [\sigma]$.

19. There are similar results to this for other notions of randomness. Notably, Eagle (2005) points out that one cannot tell from any finite initial segment of a sequence that it is *von Mises random* if it is. He concludes that von Mises randomness "is a profligate hypothesis that we cannot be justified in adopting" even for infinite strings of quantum-mechanical measurement outcomes (757–58). Eagle suggests Martin-Löf randomness as an improvement on von Mises's notion. But, given the sequence of arguments here, he would presumably take Martin-Löf randomness and Schnorr randomness to be similarly profligate, especially when, as we will see, there is a sense in which empirical underdetermination holds here even in the limit as one examines the entire sequence. In contrast, we take both of these notions to be in the ballpark of the physical intuitions involved in quantum randomness. We discuss Eagle's approach to randomness later.

https://doi.org/10.1086/708712 Published online by Cambridge University Press

The upshot is that there is no effective procedure that would tell whether a particular sequence is MLR or SR but not MLR.

In order to make clear what is at stake here, consider two ways of understanding what it might mean for the sequence of results σ in the quantum coin-toss experiment to be randomly determined dynamically.

> *Martin-Löf dynamics*: When a measurement is made of system S_k, its state instantaneously jumps to an eigenstate of the observable being measured in such a way that the sequence of results σ should be expected almost always to be MLR.

> *Schnorr dynamics*: When a measurement is made of system S_k, its state instantaneously jumps to an eigenstate of the observable being measured in such a way that the sequence of results σ should be expected almost always to be SR.

Given proposition 2, there is a sense in which these two dynamical laws are effectively indistinguishable, but the Martin-Löf dynamics is in fact more restrictive from a god's-eye view than the Schnorr dynamics. This difference would only be detectable by a computationally strong observer, one who could carry out computations that go beyond what can be accomplished by an ordinary Turing machine. But such an observer might find herself with very strong empirical evidence for accepting the Schnorr dynamics over the Martin-Löf dynamics.

Suppose that one somehow knew that the sequence of results σ was SR but not MLR. Since one would expect σ to be both SR and MLR on the Martin-Löf dynamics, this would count as very strong evidence in favor of the Schnorr dynamics over the Martin-Löf dynamics. This is precisely analogous to the argument that getting something from the measure-zero set of sequences that can be represented by finite algorithms would provide strong empirical evidence that the actual physical dynamics was not random at all.

That said, if the Schnorr dynamics were in fact descriptive of the world, while such a sequence of results is possible, one would never expect a sequence that was SR but not MLR. Rather, one would fully expect σ to be both MLR and SR on both the Martin-Löf dynamics and the Schnorr dynamics. Inasmuch as the two laws yield precisely the same expectations, there is good reason to take them to be empirically equivalent even after one concedes that it is logically possible for an observer to have evidence in favor of one over the other and a computationally strong observer to recognize the difference.

But to see why this does not settle the matter, consider the following nonstandard law.

https://doi.org/10.1086/708712 Published online by Cambridge University Press

Nonstandard dynamics: When a measurement is made of the system S_k, its state instantaneously jumps to an eigenstate of the observable being measured in such a way that the sequence of results σ should be expected almost always to be SR but not MLR.

Since one should expect this dynamics to produce a sequence of measurement outcomes that is SR, one should expect it to produce a sequence that appears to be perfectly random in the Schnorr sense of not exhibiting any effectively specifiable or discernible pattern. Among other things, this means that one should expect all initial segments of the sequence of measurement results to appear to be completely arbitrary and patternless in every algorithmically specifiable sense. But inasmuch as one should expect the full sequence to be SR but not MLR, one should expect that it will be selected from a measure-zero set of infinite-length sequences. So while this dynamics produces sequences whose initial segments will always appear to be entirely patternless and unpredictable and will pass all effective statistical tests for being random, a sequence chosen from a measure-zero set is in a straightforward sense very special and, hence, is not at all random in the measure-theoretic sense. While the sequence of measurement results will appear to be randomly determined on the nonstandard dynamics, it is not.

Similarly, while one should expect results produced by the nonstandard dynamics to appear to be statistically independent, they are not. If the results were in fact statistically independent, then the sequence should be expected to be arbitrarily chosen from the measure-one set of all possible infinite-length sequences, not from the measure-zero set of sequences that are SR but not MLR. Hence, the sequence of results produced by the nonstandard dynamics is not random in the sense of in fact being statistically independent.[20]

While the nonstandard dynamics represents a simple, concrete law that an inquirer might seriously consider given standard deliberational resources, it threatens a strong variety of empirical underdetermination. Since there is no effective procedure that would distinguish between a sequence that is both SR and MLR from one that is SR but not MLR, the nonstandard dynamics is empirically equivalent to the Martin-Löf dynamics given standard computational resources. But inasmuch as one should expect the Martin-Löf dynamics to be empirically indistinguishable from process 1, the nonstandard dynamics is empirically equivalent to process 1 if one is restricted to standard computational resources. Indeed, it is empirically equivalent to

20. If one were to repeat the full quantum coin-toss experiment and keep getting sequences in the gap between SR and MLR, then from a god's-eye view one would have evidence for a very subtle sort of global statistical dependence—a sort that one could not concretely characterize by effective means.

https://doi.org/10.1086/708712 Published online by Cambridge University Press

any standard criterion of randomness that assigns Lebesgue measure one to the set of random sequences. The upshot is that if the sequence of quantum results σ is in fact random in any standard sense, then there is no effective way to rule out the nonstandard dynamics regardless of how much empirical evidence one has.

The epistemic situation here is dire. There is a straightforward sense in which one can never have any empirical evidence whatsoever that one's quantum-mechanical results are in fact randomly determined or genuinely independent if they are. In order to see why, compare what it would be to have empirical evidence regarding the *relative frequencies* of one's results against what it would be to have empirical evidence regarding the *randomness* of one's results.

A good Bayesian inquirer might have empirical evidence either for or against quantum mechanics predicting the right relative frequencies by conditioning on the evidence presented in each initial segment σ_k of σ on the assumption that the sequence exhibits an appropriate sort of statistical uniformity. But there is no way at all to distinguish between the initial segments of sequences that are both SR and MLR and those that are SR but not MLR. This is because c-incompressible on a prefix-free machine (the condition for being MLR) is precisely the same thing as c-incompressible on a computable measure machine (the condition for being SR) for any finite initial segment of the sequence. Sequences that are both SR and MLR and those that are SR but not MLR will both always appear to be similarly random, patternless, statistically independent, and unpredictable.

Further, because the conditions are identical for all finite initial segments, no background assumption of uniformity for the full sequence will help a Bayesian inquirer to distinguish between sequences that are both SR and MLR (and hence genuinely random) and those that are SR but not MLR and hence not what one would expect from a random process. The point here is not that the inquirer will never know with certainty whether the sequence was randomly determined. Rather, even with a background assumption that the string is overall statistically uniform, looking at finite initial segments here provides no evidence whatsoever that the sequence was in fact randomly determined.

While the examination of initial segments might provide a Bayesian inquirer with compelling evidence that a given sequence is or is not simply patterned in a concrete specified way (as in the case of the alternating pattern exhibited by the sequence 01010101. . .), a sequence that is SR but not MLR has a global property that cannot be detected by examining initial segments. A sequence generated by the nonstandard dynamics should be expected to exhibit this global property, one shared by measure-zero of the possible infinite-length sequences. Such sequences are very special. But the fact that they are special is not detectable from finite initial segments.

https://doi.org/10.1086/708712 Published online by Cambridge University Press

While a good Bayesian is not committed to any particular set of priors as being rational, on the standard line at least, she is committed to probabilistic coherence and nondogmatic priors. The first condition allows her to avoid dutch books, and the second provides a general path to learning the truth. If a Bayesian inquirer were ever to assign a probability of zero or one to a hypothesis under consideration, she would never be able to condition away from the initial dogmatic assignment and hence would be entirely insensitive to new evidence no matter how strong. While she might assign a very low prior probability to the hypothesis that the evidence is SR but not MLR, inasmuch as she is interested in the truth, she cannot assign a probability of zero. But, once on the table, she would never have empirical evidence that supports both SR and MLR over SR but not MLR if she is restricted to standard computational resources.

That a good Bayesian agent may have evidence regarding limiting relative frequencies illustrates that the epistemic problem here is not a version of the standard problem of induction. Even when an agent has *full* empirical information in the form of the complete infinite-length sequence σ, she can have no empirical evidence at all regarding whether σ is both SR and MLR (and hence intuitively compatible with any standard notion of randomness) or SR but not MLR (and hence intuitively incompatible with all standard notions of randomness) given standard computational resources.

4. Discussion. The argument here concerns both the objective nature of the physical world and what we can know about it. A theory that says that a process is random in a particular sense can only be true if the process is in fact random in the sense specified. That said, the claim is not that quantum randomness is in fact faithfully described by any particular notion of algorithmic randomness that we have discussed. We can think of no reason whatsoever to suppose that quantum randomness as exhibited in the physical world is precisely characterized by SR but not MLR or by the significantly more plausible notions of Schnorr randomness or Martin-Löf randomness themselves. These notions of randomness are defined in terms of our basic understanding of formal computability, and there is no grounds for believing that physical law respects that. Rather, the argument is that objective standards of randomness that are in the ballpark of capturing the properties that we expect from the random results of the quantum coin-flip experiment, notions like SR and MLR, are subtle enough as to be computationally (and hence empirically) indistinguishable. Further, in the case of SR but not MLR, we see how one might have a notion of randomness that one expects will select sequences from a measure-zero subset of possible sequences, and hence violate a basic commitment concerning the nature of objectively random quantum processes (i.e., that they select typical sequences), yet is nevertheless computationally indistinguishable from other standards of randomness that are in the ballpark

https://doi.org/10.1086/708712 Published online by Cambridge University Press

of capturing the properties that we expect from the random measurement results. The upshot is that we cannot empirically distinguish between very different dynamical laws by standard computational means.

Among the consequences of quantum mechanics is presumably the claim that the quantum coin-toss experiment in fact models an objectively random fair coin and hence selects a typical sequence in the Lebesgue measure-one sense from the set of all possible sequences. In something like the standard collapse formulation of quantum mechanics, the truth of a dynamical law of nature is at stake—the dynamics of the world is such that either one should expect one's results to be randomly distributed in a particular specified way or they are not. More concretely, most physicists would presumably expect the actual results of a quantum coin toss to be both SR and MLR if they were to consider the question. As we have seen, both of these notions plausibly capture what it means for a sequence to be patternless and unpredictable, and they satisfy the measure-one typicality intuition regarding what it means to be a fair coin.

Of course, the point concerning the objectivity of quantum randomness also applies in the context of a deterministic hidden-variable theory like Bohmian mechanics. Here there is a physical matter of fact at stake concerning whether particles are in fact randomly distributed in a particular specified sense with respect to the wave function at a time. If they are, then the results of the quantum coin-toss experiment should be expected to be randomly distributed in a corresponding sense; otherwise, all bets are off.

One strategy for avoiding the empirical underdetermination of quantum randomness would be to appeal to a different notion of what it means for a process to be random. Eagle (2005) has argued that randomness might be understood as just unpredictability for a specified predictor. While the algorithmic notions of randomness that we have considered are grounded in the martingale idea that a random sequence should be unpredictable at a specified level of computation, his proposal is more practically minded. Specifically, Eagle takes an event E to be random for a predictor P with a theory T if and only if P's posterior probability of E conditional on T and her current evidence, is equal to her prior probability of E. The idea is that the event is random if and only if a human agent P cannot in fact make better predictions given her evidence of the current state than she could with just her theory. Eagle wants a notion of randomness that takes into account the fact that real predictors are severely limited in their epistemic and computational capacities. He takes his account of randomness to be objective in the sense that it is based on the objective features of real predictive agents.

Shifting to a notion of randomness like this would arguably allow one to have straightforward evidence regarding the randomness of quantum events relative, say, to the actual community of physicists given their de facto epistemic capabilities. The results of the quantum coin toss experiment are random

https://doi.org/10.1086/708712 Published online by Cambridge University Press

on this view if and only if the physical community can do no better than to predict the standard quantum probabilities. Since this has repeatedly proved to be the case, we have substantial empirical evidence for the randomness of quantum measurement results relative to the actual physics community over its recent history given its formulation of quantum mechanics and its epistemic access to facts about the physical world.

Such a notion of randomness makes quantum randomness a property that depends on the contingent practical properties of human agents rather than an objective intrinsic property of the physical world. But inasmuch as one is concerned with determining whether quantum mechanics has the physical world right, the question is not one whether the community of physicists up to now has been able to make predictions better than quantum mechanics allows when they use quantum mechanics; rather, it is whether quantum mechanics is right in characterizing the collapse as being a fundamentally random physical process in the context of something like the standard collapse theory or of the initial distribution of particles being genuinely random in the context of something like Bohmian mechanics. The question is whether the physical world is in fact random in a concrete sense that supports the descriptions of our best physical theories. Also salient here, the expected measure-theoretic properties of random quantum sequences seem to be an essential part of our theoretical commitments. In the case of the quantum coin-flip experiment, the question is whether the sequence of results can be expected to be typical—that is, selected in an unconstrained way from a Lebesgue measure-one subset of the set of all sequences. This reflects von Neumann's theoretical commitment to the results being *arbitrary* and *capricious*.[21]

Notions of algorithmic randomness allow us to consider ideal predictors with computational abilities that outstrip our current, contingent abilities. This allows us to specify part of the standard commitments of physicists with respect to quantum mechanics, namely, those commitments that go beyond our historically contingent capacities to predict. But, as we have seen, this comes with a trade-off—the tools of algorithmic randomness allow us to specify notions of randomness such that we cannot determine whether they in fact obtain.

5. Epistemic Morals. We have seen how one might simply characterize a set of sequences (those that are SR but not MLR) where each will always appear to be patternless and will be empirically indistinguishable from a

21. While von Neumann was committed to quantum mechanics being indeterministic, it is worth noting that algorithmic notions of randomness are compatible with determinism. As a simple example, given the right sort of initial distribution of particles, the quantum world may be deterministic as described by Bohmian mechanics and yet still yield results that are expected to be MLR.

https://doi.org/10.1086/708712 Published online by Cambridge University Press

standard measure-one notion of randomness like MLR given computable resources. If nature were always to produce sequences drawn from this measure-zero set, the quantum world would not be random, but one could never know this by effective means. Indeed, as we have seen, there is a sense in which one could never have any empirical evidence at all for accepting a standard random dynamics if one were ever to allow for the possibility of something like the nonstandard dynamics obtaining. And it is unclear the rational grounds on which one might rule out this entirely straightforward possibility—a possibility that might easily be tested if one had nonstandard computational resources of sufficient strength.

This leaves us with a sort of empirical underdetermination that results from computational limits and not from any lack of empirical evidence. Even with the complete set of empirical evidence, full Turing computational power, and the assumption that one's data are statistically uniform, there is a clear sense in which one can have no empirical justification whatsoever for believing that the results of one's quantum-mechanical experiments are arbitrarily, independently, or randomly determined.

The practice of science often involves theoretical commitments that extend beyond our predictive capacities. These might concern the energy density of the early universe or the continuity of space-time or the expected properties of infinite sequences of quantum coin flips. Such idealized commitments often help us in formulating, reasoning about, and communicating the content of our physical theories. But they do so at the cost of committing us to claims that may not be empirically testable given our actual empirical and computational capabilities.[22]

In the present case, if one is limited to computable resources, there is no empirical content to insisting that quantum-mechanical results are genuinely random, arbitrary, independent, or patternless. While one might be fully committed to their being randomly determined given one's intuitions regarding dynamical simplicity or naturalness, one can have no empirical evidence for so believing.

REFERENCES

Albert, David Z. 1992. *Quantum Mechanics and Experience*. Cambridge, MA: Harvard University Press.

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

———. 2019. *The Conceptual Foundations of Quantum Mechanics*. Oxford: Oxford University Press.

22. Different ways of modeling such idealized commitments can also matter to the empirical testability of scientific theories. This is discussed in the context of nonstandard probability theory in Huttegger (2019).

https://doi.org/10.1086/708712 Published online by Cambridge University Press

Bohm, David. 1952. "A Suggested Interpretation of Quantum Theory in Terms of 'Hidden Variables.'" Pts. 1 and 2. *Physical Review* 85:166–79, 180–93.
Dirac, P. A. M. 1957. *The Principles of Quantum Mechanics*. Oxford: Oxford University Press.
Downey, Rodney G., and Evan J. Griffiths. 2002. "Schnorr Randomness." *Electronic Notes in Theoretical Computer Science* 66 (1): 25–35.
Downey, Rodney G., and Denis R. Hirschfeldt. 2010. *Algorithmic Randomness and Complexity*. Berlin: Springer.
Eagle, Anthony. 2005. "Randomness Is Unpredictability." *British Journal for the Philosophy of Science* 56:749–90.
Earman, John. 1986. *A Primer on Determinism*. Dordrecht: Reidel.
Ghirardi, G. C., A. Rimini, and T. Weber. 1986. "Unified Dynamics for Microscopic and Macroscopic Systems." *Physical Review* D 34:470–91.
Huttegger, Simon M. 2019. "Rethinking Convergence to the Truth." Unpublished manuscript, University of California, Irvine.
Li, Ming, and Paul Vitányi. 2008. *An Introduction to Kolmogorov Complexity and Its Applications*. 3rd ed. New York: Springer.
Saunders, Simon, Jonathan Barrett, Adrian Kent, and David Wallace, eds. 2010. *Many Worlds? Everett, Quantum Theory, and Reality*. Oxford: Oxford University Press.
Shen, A., V. A. Uspensky, and N. Vereshchagin. 2017. *Kolmogorov Complexity and Algorithmic Randomness*. Mathematical Surveys and Monographs 220. Providence, RI: American Mathematical Society.
Soare, Robert. 2016. *Turing Computability: Theory and Applications*. Berlin: Springer.
von Neumann, J. 1955. *Mathematical Foundations of Quantum Mechanics*. Princeton, NJ: Princeton University Press. Trans. R. Beyer from *Mathematische Grundlagen der Quantenmechanik* (Berlin: Springer, 1932).
Wallace, David. 2012. *The Emergent Multiverse: Quantum Theory according to the Everett Interpretation*. Oxford: Oxford University Press.
Wigner, Eugene P. 1970. "On Hidden Variables and Quantum Mechanical Probabilities." *American Journal of Physics* 39:1005–9.

https://doi.org/10.1086/708712 Published online by Cambridge University Press