THE PRINCIPLE OF SUFFICIENT REASON AND PROBABILITY

ALEXANDER R. PRUSS

Abstract. I shall argue that considerations about frequency-to-chance inferences make very plausible some a localized version of the Principle of Sufficient Reason (PSR). But a localized version isn't enough, and so we should accept a full PSR.

1. The Principle of Sufficient Reason

I will take the Principle of Sufficient Reason (PSR) to be the claim that necessarily every contingent truth has an explanation (Pruss 2006).

It has been often noted that many of our ordinary epistemic practices presuppose the existence of explanations. Rescher, for instance, talks of how the investigators of an airplane crash do not conclude that there is no explanation from their inability to find an explanation (Rescher 1995, p. 2). But it is one thing to agree that the PSR holds for propositions about everyday localized matters, and another to generalize to cosmic cases, such as the origination of the universe. (Of course it would be difficult to precisely define what counts as "localized".) Still, we have a rough grasp of localization: an LED lighting up is localized, while an infinite regress of events or the complete state of our large universe presumably are not.

I will argue, however, that in order to make sense of our scientific epistemic practices, we need a principle like the PSR that applies to global matters. In doing so, I will draw on recent mathematical work on laws of large numbers for nonmeasurable processes (Pruss 2013).

The argument begins with the local cases and generalizes to the global ones.

I need to note that I understand the PSR to be compatible with indeterministic phenomena. The PSR is a claim that explanations exist, not that deterministic explanations exist. The explanations might well be probabilistic in nature, and might even involve low probabilities, as is widely accepted in the philosophy of science (Salmon 1989). After all, giving explanations is largely about making events understandable, and we understand low probability chancy events as well as high probability ones (Jeffrey 1969).

Forthcoming in *Oxford Studies in Metaphysics*.

2. Local chances

2.1. **Frequency to chance inferences.** I will consistently use the word "chance" to indicate probabilities that are the tendencies of stochastic processes. Consider now this simple inference. A coin has been independently flipped a thousand times, and about 750 times the coin has landed heads. There is a very natural inference: The coin is loaded in such a way as to have approximately a $3/4$ chance of landing heads. This natural inference leads to further predictions of the coin's behavior.

Frequency-to-chance inferences like the above are everywhere, and inductive reasoning to a universal generalization is arguably just a limiting case.

Let C_p be the hypothesis that the chance of the coin's landing heads is p. Plausibly, the inference in the coin example was based on the fact that given C_p with p close to $3/4$, it is very likely that about 750 times we will get heads, while given C_p with p far from $3/4$, this is unlikely. The reason for this is a Law of Large Numbers: given a large number of independent and identically distributed chancy trials, the frequency of an outcome among the trials is likely to be close to the chance of the outcome.

There are difficulties here, of course, with how we determine that the coin flips are independent and identically distributed. Independence perhaps is backed by the fact that we just cannot find any memory mechanism for the coin, and identical distribution by the fact that the coin flips appear to be identical in all relevant respects. The questions of how exactly one cashes out these considerations are difficult, but they are not the questions I want to focus on.

Instead, I want to focus on why there are probabilities at all. Our intuitive Bayesian-flavored reasoning was based on comparing the probability of our observed frequencies on the different hypotheses C_p. But what about hypotheses on which there are no probabilities of frequencies? In the next subsection, I will sketch a picture of how such hypotheses might look by making use of nonmeasurable sets. A reader who wishes to avoid mathematical complication may wish to skip that subsection.

The basic idea behind the technicalities is to imagine a case where we pick out a series of heads and tails results by throwing a dart at a circular target, and deeming heads to have occurred when the dart lands within some set A and tails to occur otherwise. But we will take the set to be utterly nonmeasurable ("saturated nonmeasurable"): it has no area, and doesn't even have a range of areas. I will then argue on technical grounds that the hypothesis that coin toss results were generated in such a way can neither be confirmed nor disconfirmed by observation. But in order to confirm an alternate hypothesis—say that a series of heads/tails results was the result of a sequence of fair coin tosses—we need to rule out the hypothesis that it was generated in some "nonemeasurable" way like this.

2.2. **Nonmeasurable sets.** So suppose coins are being flipped in this roundabout way. A dart with a perfectly defined tip (say, infinitely sharp or perfectly symmetrical) is uniformly randomly thrown at a circular target. There is a region A of the target marked off, and a detector generates a heads toss (maybe in a very physical way: it picks up the coin and places it heads up) whenever the dart lands in A; otherwise, it generates a tails toss.

The chance of the coin landing heads now should be equal to the proportion of the area of the target lying within A, and we can elaborate C_p to the hypothesis that the area of A is pT, where T is the area of the whole target. Given the observation of approximately 750 heads, it seems reasonable to infer that probably the area of A is approximately $(3/4)T$.

But what if an area cannot be assigned to the marked region A? Famously, given the Axiom of Choice, there are sets that have no area—not in the sense that they have zero area (like the empty set, a singleton or a line-segment), but in the sense that our standard Lebesgue area measure cannot assign them any area, not even zero. In such a case, there will also be no well-defined[1] chance of the dart hitting A, and hence no well-defined chance of heads. Let N be the hypothesis that A has no area, i.e., is non-measurable.

Now, here is a fascinating question. If N were true, what would we be likely to observe? Of course, if we perform 1000 tries, we will get *some* number n of heads, and $n/1000$ will then be a frequency between 0 and 1. We might now think as follows. This frequency can equally well be any of the 1001 different numbers in the sequence $0.000, 0.001, \ldots, 0.999, 1.000$. It seems unlikely, then, that it's going to be near 0.750, and so $C_{3/4}$ (and its neighbors) is still the best hypothesis given that the actual frequency is 0.750.

But this reasoning is mistaken. To see this, we need to sharpen our hypothesis N a little more. There are non-measurable sets A where we can say things about the frequency with which A will be hit. For instance, it could be that although A is non-measurable, it contains a measurable set A_1 of area $(0.74)T$ and is contained in a measurable set A_2 of area $(0.76)T$. (Think of A as 74% of the target plus a nonmeasurable set localized to an area containing only 2% of the target.) But the dart will hit A_1 about 74% of the time, and A_2 about 76% of the time. Whenever the dart hits A_1, it hits A, and whenever it hits A, it hits A_2, so we would expect the dart to hit A approximately 74% to 76% of the time. And so our observed frequency would be no surprise.

The mere fact that A is nonmeasurable does not rule out probabilistic predictions about frequencies, because a nonmeasurable set might be "quite close" to measurable sets like A_1 and A_2 that bracket A from below and above.

However some sets are not only nonmeasurable, but *saturated nonmeasurable*. A set A is saturated nonmeasurable provided that all of A's measurable subsets have measure zero and all measurable subsets of the complement of

[1] One wants to say "no chance", but that suggests "zero chance", which is not what one is saying.

A are also of measure zero. Given the Axiom of Choice, for Lebesgue measure on the real line, there not only are nonmeasurable sets, but there are saturated nonmeasurable sets (Halperin 1951).[2]

When A is saturated nonmeasurable, no method of generating predictions by bracketing a the probabilities into an interval, like the one from 74% to 76%, will work. The only measurable subsets of A will have zero area and the only measurable supersets of A will have area T. So our bracketing will only tell us the trivial fact that the frequency will be between 0 and 1, inclusive.

Let M then be the hypothesis that A is saturated nonmeasurable. Can we say that given M, the frequency is unlikely to be near 0.750?

The answer turns out to be negative even if we have *infinitely* many observations. Pruss (2013) has given a plausible mathematical model of an infinite sequence of independent identically distributed saturated nonmeasurable random variables, and it follows from his Theorem 1.3 that for any nonempty interval I which is a proper subset of $[0, 1]$, the event that the limiting frequency of the events is in I is itself saturated nonmeasurable. Thus, even with infinitely many independent shots, we could say nothing probabilistic about the observed frequency of hits of our saturated nonmeasurable set being near 0.750: the probability would neither be large nor small. And analogous claims can be argued for in finite cases by building on the tools in that paper (see Theorem 2 in the Appendix).

So in our case above, if I is some small interval like $[0.740, 0.760]$ centered on 0.750, there is nothing probabilistic we can say about the observed frequency being in I when the target set is nonmeasurable as the event of the frequency being in I is then saturated nonmeasurable. In particular, we cannot say that the frequency is unlikely to be near 0.750 (nor that it's likely to be near). The observation of a frequency close to 0.750 is neither surprising nor to be expected given that the target set is saturated nonmeasurable.

Our probabilistic reasoning thus cannot disconfirm hypotheses of saturated nonmeasurability. Such hypotheses endanger all our local scientific inferences from observed frequencies to chancy dispositions, inferences central to our epistemic practices. Yet our local scientific inferences are, surely, good. If we cannot disconfirm saturated nonmeasurability hypotheses *a posteriori*, then we need to do so *a priori*.

Kleinschmidt (2013) defends a presumption of explanation principle ("EP") on which we are justified in presuming explanation, by analogy with our presumption that a student coming late to class was not struck by lightning, even though people are struct by lightning. She then writes: "In the lightning case, we can explain our reluctance by noting that we justifiably

[2] If A_0 is a saturated nonmeasurable subset of the interval $[0, 1]$, then the Cartesian product $A_0 \times [0, 1]$ will be a saturated nonmeasurable subset of the square $[0, 1]^2$. (There is also a directly two-dimensional construction in Sierpiński 1938.) Intersecting this with a disc of radius $1/2$ centered on $(1/2, 1/2)$ yields a saturated nonmeasurable region for a disc-shaped area, like in our example.

believe, partly due to induction, that lightning strikes rarely. We might hope to give a similar explanation of our endorsement of EP." But while we can say what it would look like if lightning was more common than we think it to be—and thus we can say on empirical grounds that it is not so—the upshot of our argument is that we cannot say what it would look like if there were more unexplained events than we think, and hence cannot say on empirical grounds that it is not so.[3]

Another approach would be to assume low *prior* credences for hypotheses of saturated nonmeasurability. But it is difficult to justify this in a way that is not *ad hoc*. Given the Axiom of Choice, there are just as many saturated nonmeasurable subsets of, say, $[0, 1]$ as there are measurable ones (see the Corollary in the Appendix). Granted, subjective Bayesians may not worry too much about the *ad hoc* in the choice of priors, but the reason they do not worry is because of the hope that evidence will swamp the priors and make them irrelevant. But evidence cannot affect the credence of the nonmeasurability hypotheses, since those hypotheses do not generate the kinds of likelihoods that are needed for confirmation or disconfirmation, and so one would simply be stuck with the *ad hoc* priors.

Note, however, that the credences of such hypotheses would have to be very low indeed: they would have to be zero or infinitesimal. I sketch the argument in a simple case. Suppose that N is a hypothesis according to which the coin tosses are independent (in the sense of Pruss 2013) saturated nonmeasurable sets. Suppose that $P(N) > 0$ and that the only alternative to N is the hypothesis H that the coin tosses are fair and independent so that $P(H) = 1 - P(N)$. Let E be the observed evidence—a sequence of n coin toss results (say $HTHHHTH$ if $n = 7$). We would like to be able to use Bayes' Theorem to say that the probability of H given the evidence is:

$$P(H \mid E) = \frac{P(E \mid H)P(H)}{P(E \mid H)P(H) + P(E \mid N)P(N)}.$$

[3]Kleinschmidt (2013) also attempts to give counterexamples to the PSR on the basis of cases involving the diachronic identity and synchronic composition of persons. Why, for instance, does cloud of atoms in group G_1 compose a person while the cloud of atoms in group G_2, which differs only by a single atom, does not? (The fission case can be handled analogously. A full response to her intricate arguments would take us too far a field, but the facts underwriting the composition claims are either necessary or contingent. Since I am only defending a PSR for contingent truths, it is only the conitngent option that I need consider. But even without assuming the PSR, simply using Kleinschmidt's own EP, it will count very strongly against a theory on which it is contingent that an arrangement of atoms composes a person if the theory can give no explanation of why it happens to do so. Consider, for instance, dualist theories on which souls attach to some groups of atoms with no explanation as to why they attach to the ones they attach to. Those theories are implausible precisely because of the explanatory failure. And the same will be true of those materialist theories on which it is a brute unexplained fact that some atoms compose a person. Kleinschmidt has not shown that there is a brute-composition theory that is still plausible after taking into account her own EP.

Unfortunately, the term $P(E \mid N)$ in the denominator is not a number, since E is saturated nonmeasurable conditionally on N. But we can, nonetheless, use Bayes' Theorem. We can model the probability of a nonmeasurable set by an interval of numbers. In the case of $P(E \mid N)$, the interval will range all the way from 0 to 1, i.e., will be $[0,1]$. Fortunately, only one of the quantities on the right-hand-side corresponds to an interval: the other quantities, $P(E \mid H)$, $P(H)$ and $P(N)$ are ordinary classical probabilities. We can then use the range of the interval $P(E \mid N)$ and Bayes' Theorem to define an interval-valued probability for $P(H \mid E)$.

Since $P(E \mid N)$ is the interval $[0,1]$, one endpoint for the interval corresponding to $P(H \mid E)$ will be given by replacing $P(E \mid N)$ in Bayes' Theorem with zero:

$$\frac{P(E \mid H)P(H)}{P(E \mid H)P(H) + (0) \cdot P(N)} = 1.$$

This turns out to be the upper endpoint. The second endpoint—the lower one—will correspond to replacing $P(E \mid N)$ with one, the other endpoint of *its* interval:

$$\frac{P(E \mid H)P(H)}{P(E \mid H)P(H) + (1) \cdot P(N)}.$$

Now, $P(E \mid H) = (1/2)^n$. It follows that this lower endpoint of the interval given by $P(E \mid H)$ is less than or equal to:$(1/2)^n/P(N)$. Therefore, the interval-valued probability $P(H \mid E)$ will contain the interval $[(1/2)^n/P(N), 1]$.[4]

As long as $P(N)$ is strictly bigger than zero and not infinitesimal, it follows that as n goes to infinity—i.e., as the amount of evidence gathered increases—the hypothesis H gets to have a probability range closer and closer to $[0,1]$. But in order to get confirmed, to become believable, H would have to be an interval of probability like $[a,b]$ where a is close to one, not close to zero. So the more evidence we gather, the broader our probability interval for H. This is Bayesian divergence rather than convergence. As we observe more coin-flips, the nasty hypothesis N infects our probabilities more and more.

One could try to claim that $P(N)$ is infinitesimal, but that would seem to be *ad hoc*. Alternately, one might try to completely *a priori* rule out pathological hypotheses, ruling them to have probability zero. For instance, one might simply deny any version of the Axiom of Choice strong enough to imply the existence of nonmeasurable sets. Or one might deny the possibility

[4] This argument is simpler than the usual kind of reasoning with interval-valued probabilities (e.g., Whitcomb 2005), because only one of the quantities involved is an interval. We can justify the argument more rigorously by thinking of interval valued probabilities as determining a family of classical probability functions, and then applying Bayes' Theorem to each member of the family. The only relevant variation between members of the family will be as to the value assigned to E given N, and that will range between 0 and 1, thereby generating a range for the posterior probability containing $[(1/2)^n/P(N), 1]$.

of physical processes sensitive to whether an input falls into a pathological set.

2.3. **No-explanation hypotheses.** But clever mathematical constructions are not the only potential source of probabilistic nonmeasurability in the physical world. Suppose a coin is tossed, but there is a hypothesis H that there will be no explanation at all of how the coin lands. Simplify by assuming that the coin must land heads or tails. Then I contend that the event of the coin landing heads has no probability conditionally on H.

For suppose it has some probability. What is it? The only non-arbitrary answer seems to be the one given by symmetry or the Principle of Indifference: 1/2. And, generalizing, the only non-arbitrary probability for any particular fixed sequence of n tosses on a no-explanation hypothesis will be 2^{-n}. But if that is the answer, then the no-explanation hypothesis cannot be empirically distinguished from the hypothesis that the coin tosses form an independent sequence of stochastic events each with probability 1/2 of heads, as both hypotheses predict the same behavior.

In fact, although 1/2 is the only non-arbitrary answer, the arbitrary answers face exactly the same problem. For each possible answer, arbitrary or not, to the question of the probability that the coin lands heads conditionally on the no-explanation hypothesis, there is a possible stochastic process hypothesizing which yields the same prediction as the no-explanation hypothesis does. And then no amount of observation can rule out the no-explanation hypothesis. Whether this is a serious problem will depend on whether we can rule out the hypothesis *a priori*, an option we will discuss later in this section.

Formal results on nonmeasurable sets like those discussed in Section 2.2 model not only those events that are saturated nonmeasurable due to their pathological set theoretic construction but also those that are nonmeasurable simply because no probabilities can be assigned to them.

Within a probabilistic framework, there are now three possibilities. Either (a) conditioning on a no-explanation hypothesis yields an exact probability, or (b) it yields no probability at all, or (c) it yields a probability range or interval. If it yields an exact probability, then we cannot empirically distinguish the no-explanation hypothesis from an explanatory chance hypothesis whose chance assignment matches that exact probability. If it yields no probability at all, not even a range of probabilities, we are dealing with a saturated nonmeasurable event, and our formal results tell us that we cannot make any probabilistic predictions about frequencies, and in particular cannot say whether it is likelier or less likely that we would get the same observations from an explanatory chance hypothesis.

Now consider the third option, that of a probability range. Perhaps when we condition on the no-explanation hypothesis, we get an interval $[a, b]$ of probabilities with $a < b$ and $[a, b] \neq [0, 1]$. In such a case we can make some probabilistic predictions by the Law of Large Numbers (Pruss 2013, inequality (1)), namely that limiting frequencies are almost surely going to be between a and b (both inclusive). It can be shown that nothing stronger can be shown about the limiting frequencies than that they are between a and b (Pruss 2013, Theorem 1.3). But a no-explanation hypothesis assigns at best the full range $[0, 1]$ to the probability, and so the only thing that follows from the no-explanation hypothesis is that the frequencies range between 0 and 1, which is trivially true. We cannot, thus, say that a particular set of observed frequencies fits worse or better with the no-explanation hypothesis than with some explanatory hypotheses: all frequencies fit equally well with the no-explanation hypothesis.

Thus, whether or not the no-explanation hypothesis has a probability, and whether this probability is exact or an interval, the no-explanation hypothesis cannot be ruled out empirically. It is simply not amenable to statistical methods, or if it is amenable, it is amenable in the wrong way: The argument based on an interval-based interpretation of Bayes' Theorem given at the end of Section 2.2 shows that once we have a non-zero non-infinitesimal probability of a no-explanation hypothesis, we end up with poorer and poorer confirmation of serious explanatory hypotheses the more data we gather.

So the no-explanation hypothesis must be ruled out *a priori*. There are two initially attractive proposals for doing so. We can say that no-explanation hypotheses simply cannot hold. That would be to subscribe to a localized version of the PSR. Or we could simply say that no-explanation hypotheses have very low prior probabilities.

But as in the case of saturated nonmeasurability, it is unsatisfactory to set low priors for no-explanation hypothesis. Any values for these priors would be arbitrary once we allowed for the possibility of their truth, and the usual subjective Bayesian hope that arbitrary priors will eventually be swamped by evidence (a classic statement of this is in Edwards, Lindman and Savage 1963, p. 197) is not applicable given the non-amenability of no-explanation hypotheses to empirical study.

One might compare the low priors for no-explanation hypotheses to the low priors we assign some other wacky hypotheses. For instance, our prior for the aberrant hypothesis that the gravitational force between two objects of mass m_1 and m_2 at distance r apart has magnitude $Gm_1m_2/r^{2+10^{-100}}$ will presumably be very small, since otherwise we could never conclude that the actual magnitude of the force is Gm_1m_2/r^2, as the aberrant hypothesis cannot be empirically distinguished from the Newtonian hypothesis.

On its face, the analogy fails because we reject the $Gm_1m_2/r^{2+10^{-100}}$ hypothesis on the grounds that it does not have the kind of simplicity we expect in explanatory hypotheses, while bald no-explanation hypotheses are actually simpler than any explanatory hypotheses.

But perhaps the analogy is this: in both the no-explanation and wacky gravity cases, our only reason for assigning low priors is that if we assigned high priors, we couldn't do science. We cannot justify the assignment of low priors in any other way, and do not need to.

Neither case of assignment of low priors without further justification is satisfactory, but it is better to be in one unsatisfactory situation than two. Furthermore, one could hope that an inductive logic of a Carnapian variety might eventually non-arbitrarily justify low probability assignments to more complex hypotheses—but this won't help with, and in fact will exacerbate, the problems with very simple no-explanation hypotheses.

Finally, the defender of a global PSR has some hope of being in *neither* unsatisfactory situation. For not only has she ruled out all no-explanation hypotheses, but she has some hope of arguing to a global explanation that assigns lower probabilities to more complex scientific hypotheses. For it may be that the global explanation is one that, like the theistic and optimalist/axiarchist global explanation proposals, makes better states of affairs more likely.[5] And simpler laws are better, both aesthetically and instrumentally.

Thus far I have argued that we need some kind of localized PSR as a metaphysical principle to rule out no-explanation hypotheses on which our observations simply have no probability. A central claim is that if nothing can be said about the probabilities of individual events, then nothing can be said about the probabilities of observing frequencies of events in a sequence, even an infinite sequence, of observations. Thus, the no-explanation hypotheses cannot be ruled out observationally, since to do that we would have to say that our observations are improbable on the no-explanation hypothesis.

A tempting response to the arguments of this section is to say that just as it counts *a posteriori* against a hypothesis when our observations have low probability on it but have high probability on an alternative, it counts *a posteriori* against a hypothesis when our observations have *no* probability on the hypothesis but have high probability on an alternative. But if *no* possible sequence of observations has any meaningful probability on a no-explanation hypothesis, then by this principle *every* possible sequence of observations would disfavor our no-explanation hypothesis (since every sequence is probable on *some* explanatory hypothesis). But if E is evidence against a hypothesis, then the negation of E will be some evidence for the hypothesis. But if evidence has no probability on the no-explanation hypothesis, the negation of that evidence also has no probability on the no-explanation hypothesis, and so by the suggested principle both E and

[5]Optimalism/axiarchism as formulated by Leslie (1979) and Rescher (2010) make the best option be necessary. But an interesting variation would be a theory on which value simply makes a world more likely.

its negation would count as evidence against the no-explanation hypothesis, which appears absurd unless the no-explanation hypothesis is impossible.[6]

3. Globalizing

3.1. **Small worlds.** Can we have a local PSR without a global one? One can try to give restricted PSR-like principles that can be applied to local states of affairs without being applicable to global ones (for a critical discussion, see Koons 2006). This would allow one to secure the advantages of the PSR for scientific reasoning while avoiding the disadvantages of (a) significant ontological commitment—after all, the PSR has been used to argue for grand hypotheses like the existence of God—and (b) van Inwagen style arguments against a global PSR (van Inwagen 1983, pp. 202–204).

It is not clear, however, whether a localized PSR avoids the problems, if it's taken to be a metaphysical principle and hence metaphysically necessary.

First of all, a localized PSR applied to a small world can produce the same kinds of ontological commitment as a global PSR applied to a big world. Plausibly, it is possible to have a walnut world—a world with a finite past whose contingent components have always consisted of a single basically unchanging but otherwise ordinary walnut[7] and its parts (cf. the ball thought experiment in Taylor 1974, pp. 105–106). When we apply the localized PSR to the walnut—which is as local as something can be—we conclude that there is an explanation of the walnut. The existence of our walnut is not explained by the activity of its parts, so this walnut's explanation would have to go beyond it. But as there is nothing contingent outside the walnut in the walnut world, the explanation would have to invoke something that exists necessarily or some kind of necessary metaphysical priniciple. And this being or principle that is necessary at the walnut world would also be necessary at the actual world by the axiom S5 of the most popular system of modal logic (S5 says that if a proposition is possibly necessary, it is necessary; necessities cannot vary between possible worlds), so we have not avoided significant ontological commitment by supposing a localized PSR. Granted, this argument uses S5, but having to deny S5 would also be a case of significant metaphysical commitment. So a localized PSR has similar commitment-engendering features to its grander relatives.

Van Inwagen's argument against the PSR is based on forming the conjunction of all contingent truths and asking for the explanation of this conjunction. If this explanation is itself a contingent truth, then we absurdly have a case of self-explanation, since in explaining the conjunction of all contingent truths, this contingent truth explains itself, being one of the conjuncts. If, on the other hand, the explanation is a necessary truth, then

[6] If a hypothesis is impossible, we may be able to give an argument for its impossibility starting with E and another argument for its impossibility starting with the negation of E.

[7] Or, more precisely, an exact duplicate of an ordinary walnut. It may be an essential property of a walnut *per se* that it come from a walnut tree.

we have a puzzle about how a necessary truth can explain a contingent one. One way to set up the puzzle is by invoking the principle that if p explains q, then p entails q. Van Inwagen explicitly affirms this principle (van Inwagen 1983, p. 203), though he uses the "sufficient reason" in place of "explanation". But in the case of explanation, this principle is widely denied in the philosophy of science in the light of the existence of stochastic explanations (see, for instance, Salmon 1989), so it is better to leave the argument as simply based on the bare intuition that a necessary truth cannot explain a contingent one, which intuition I will call the "Van Inwagen Principle" (VIP).

But our best story about how the walnut in a walnut world can be explained involved the explanation of the walnut's existence by means of a necessary being or a necessary metaphysical principle. The principle-based explanation could perhaps involve something like a stochastic axiarchism, where the chance of a universe arising is proportional to its value, but whatever that would look like, it would be a violation of VIP. And the most obvious way to take the explanation of the walnut's existence by means of a necessary being would be to take it that necessary truths about the existence and essential properties of the being explain why the walnut exists, in violation of VIP.[8]

So as long as the localized PSR is a metaphysical principle, we do not actually gain much if anything from eschewing a global principle, since we get results similar to those of a global PSR when we apply the localized PSR in a small world.

Dean Zimmerman has suggested to me that locality could also be taken relative to the size of the whole universe. On the reading, a localized PSR would apply to, say, the molecules of the walnut in the walnut world, but not to the walnut as a whole. But it would be implausible to think that a walnut does not need an explanation when it is all alone, so that its existence is not localized, but needs one when it is but one among many things.

3.2. **Global probabilities.** But there is further reason to accept a global PSR. A localized PSR is insufficient to solve the kinds of probabilistic problems that led us to it.

[8]Someone who resists a global PSR but accepts a local one might also try to explain the walnut's existence by supposing an infinite regress of contingent properties of a necessary being. Perhaps the necessary being is a person who decided at t_0 to create the walnut, and did so on the basis of a contingent mental state that was caused by the person's contingent thoughts at t_{-1}; these thoughts were then caused by yet earlier contingent thoughts at t_{-2}, and so on. In this case, a localized PSR is satisfied: each localized event—the existence of the walnut, the choice at t_0, the thoughts at t_{-1}, and so on—is explained. The whole sequence of thoughts is not explained, however, which blocks the van Inwagen argument. Notice, however, that we only avoid violation of VIP by further ontological commitment, namely supposing an infinite regress of contingent states of a necessary being. Furthermore, this ontological commitment runs afoul of the modal intuition that there is a possible world containing the walnut and only at most finitely many contingent states that aren't states of the walnut.

Consider universes which have a finite past but exist over an interval of times open at the lower end, say an interval modeled by the semi-infinite interval $\{t : t > 0\}$. For each instantaneous state of such a universe, there is an earlier state, but there is no *earliest* state. It has been argued by Grünbaum (1993) that our universe is like that. Let us further assume that each state is explained by an earlier state.

A localized PSR will not apply to the infinite regresses of states we find in such universes. Thus, a localized PSR will be compatible with an unexplained no-beginning finite-past universe.[9] But then there will be no meaningful chances for the existence of such universes, and hence no non-arbitrary reason to prefer the hypothesis we think is true—a universe that begins with a Big Bang about 13.8 billion years ago—over aberrant hypotheses such as Russell's five-minute hypothesis on which the universe is five minutes old with the twist that there is no initial moment (compare Koons 2006, Section 2.1).

This was too quick. Although there are no chances associated with the two hypotheses, maybe there are epistemic probabilities. But what reason do we have to take the precise limiting sequence of states that we think obtained after the Big Bang to be more probable than the limiting sequence of states in the five minute hypothesis? Again, we would have to assign an arbitrary epistemic probability. One might think that of all possible unexplained no-beginning finite-past universes, the five-minute universe (i.e., the one that coincides with what we take to be our universe over the past five minutes, without an initial time) has an initial segment that is too elaborate—festooned with galaxies and stars as it is, and with at least one planet of very complex life—to come into existence without an explanation, while the Big Bang universe's initial segment is not as elaborate. But both initial sequences of states live in the same configuration space, and there does not appear to be a reasonable way to privilege the Big Bang sequence over the five-minutes-ago sequence as an option for coming into existence causelessly *ex nihilo*. For what it's worth there is even an intuition preferring the five-minutes-ago sequence. For one would intuitively expect an unexplained universe would have high entropy, and the states in the five-minutes-ago sequence have higher entropy than the states in the post-Big-Bang sequence do.

But suppose that we have some way of arguing that the Big Bang hypothesis is simpler than the five-minute hypothesis. We can beat the Big Bang hypothesis in respect of simplicity. Consider the hypothesis that reality consists of short truncated backwards light cone (say, extending backwards one minute), with no initial state, centered on the present state of my brain. This hypothesis supposes only a tiny fraction of the complexity of our gigantic universe (indeed even of our solar system), and there will be a lot

[9]Hume (1779, Part IX) has argued that an infinite causal regress is explained by the causes within it. However the discussion of the beginningless cannonball flight thought experiment in Pruss (2006, Section 3.1.4) makes this implausible.

more simplicity than in the Big Bang story. After all, it is a part of the Big Bang story that the world evolved in the kinds of complex ways that resulted in a life full of earth, many galaxies, and so on. All the details of this complexity will have free parameters, either determined at the time of the Big Bang or set by later chancy events.

Thus *a priori* simplicity arguments do not rule out our aberrant hypotheses. We really do seem to need a PSR.

4. PSR-compatible aberrant hypotheses

Our argument for the PSR was that it allows us to reject aberrant local and global hypotheses that we have no *a posteriori* way of rejecting. A natural response is that there are aberrant hypotheses that are compatible with the PSR as well, and the PSR only pushes the problem back a step.

For consider the hypothesis that a necessarily existing being with the essential property of being a science-hater creates a universe that appears very much other than how it is. If the PSR is restricted to contingent states of affairs, as in Pruss (2006), then it is compatible with this aberrant hypothesis. Compare this to the hypothesis that a necessarily existing being with the essential property of being a science-lover creates a universe that allows us to make correct scientific predictions. We cannot compare the science-lover and science-hater hypotheses by means of chances. By S5, each hypothesis is either necessarily true or necessarily false, and hence has chance 1 or chance 0, but we don't have access to these chances except by figuring out which hypothesis is true.

However, in the realm of PSR-compatible hypotheses, we can make use of simplicity reasoning to rule out aberrance. For instance, the science-hater hypothesis has many free parameters: just how much does the science-hater hate science and in what respects, how powerful is it, etc. And there does not seem to be any canonical set of values for these free parameters. On the other hand, there is a canonical science-lover hypothesis in the literature, namely the theistic hypothesis that there exists a perfect being, who has maximal power, maximal knowledge and loves all that is fine and good precisely in proportion to how fine and good it is. And of course science getting at the truth *is* fine and good.

And while there will be many variants of the stochastic axiarchic hypothesis that universes are likely to come into existence to the extent that they are good, depending on just how the direct relationship between value and chance works, it could turn out that a broad range of them would produce non-aberrant stories.

5. Conclusions

Where there are no explanations, there are no chances. Where there are no chances, it is difficult to have non-arbitrary probabilities. In some cases, we can have non-arbitrary probabilities based on simplicity considerations.

In others, perhaps by indifference considerations. But neither simplicity nor indifference considerations help with rejecting aberrant no-explanation hypotheses that undercut the scientific enterprise. We need a PSR, and not just one restricted to local situations, but a global one.

Finally, while our arguments above took for granted that chances were governed by the axioms of classical probability theory, it may be that we should relax this assumption. For instance, while a rational agent tends to act in proportion to the strength of her reasons, it would be hasty suppose that this tendency can be quantified by classical probability theory. Moreover, global explanatory hypotheses are unlikely to give rise to numerical chances, if only for cardinality reasons—likely, there is no *set* of all possible worlds (e.g., Pruss 2001), and classical probability theory is defined over sets of possibilities. But the main point of the paper should remain: where there are no explanations, there are no chances—whether the chances are understood as classical probabilities or are governed by some more complex calculus—and hypotheses that do not give rise to some sort of chances need *a priori* refutation.

Appendix: Some results about nonmeasurable sets

A set is saturated nonmeasurable with respect to a probability measure provided that all measurable subsets have measure zero and all measurable supersets have measure one.

Theorem 1. *Assuming the Axiom of Choice, the cardinality of the set of all saturated nonmeasurable subsets of* $[0, 1)$ *is equal to the cardinality of the set of subsets of* $[0, 1)$*.*

For the proof, we need the following easy Lemma. Let $A \Delta B$ be the symmetric difference $(A - B) \cup (B - A)$ of the sets A and B. Recall that every subset of a set of Lebesgue measure zero is measurable and has measure zero.

Lemma 1. *If* A *is a saturated nonmeasurable subset of* $[0, 1)$ *and* $A \Delta B$ *is a set of Lebesgue measure zero, then* B *is saturated nonmeasurable.*

Proof of Lemma 1. For any set C, the set $B \cap C$ differs from $A \cap C$ at most by a set of measure zero since $A \Delta B$ has measure zero, and any two sets that differ by a set of measure zero must either both be measurable or both nonmeasurable.

Thus, if C is a measurable subset of B, it follows that $A \cap C$ differs from C by a set of at most measure zero. Thus, $A \cap C$ is a measurable subset of A, hence has measure zero, and hence C also has measure zero since it differs from it by at most a set of measure zero. This argument shows that if all measurable subsets of A have measure zero, the same is true for B. Applying this argument to the complements $[0, 1) - A$ and $[0, 1) - B$, and using the observation that all measurable supersets of U that are subsets of $[0, 1)$ have measure 1 if and only if all measurable subsets of $[0, 1) - U$ have measure zero, we conclude that all measurable subset of $[0,1)-B$ have measure zero, from which we conclude that all measurable supersets of B that are subsets of $[0,1)$ have measure 1. Thus, B is saturated nonmeasurable. $\square$

Proof of Theorem 1. Famously, the Cantor middle-thirds set $C \subseteq [0,1)$ is a set of zero measure that has the same cardinality $\mathfrak{c}$ as all of $[0,1)$. Given any one saturated nonmeasurable subset S of $[0,1)$ (these exist by Halperin 1951), for each subset A of C let $S_A = (S - A) \cup A$. Then $S_A \,\Delta\, S \subseteq C$, and hence S_A is saturated nonmeasurable. There are as many sets S_A as subsets of C, and there are as many subsets of C as of $[0,1)$. $\square$

The fact that the number of subsets of the Cantor set equals the number of subsets of $[0,1)$ and that every subset of a set of zero measure is measurable also yields:

Corollary. *Assuming the Axiom of Choice, the cardinality of the set of saturated nonmeasurable sets is equal to the cardinality of the set of measurable subsets of* $[0,1)$.

Following Pruss (2013), I will model independent sequences of nonmeasurable events $E_1,\dots,E_n$ as follows. We have a probability space $(\Omega,\mathscr{F},P)$ which is a product of probability spaces $\langle\Omega_k,\mathscr{F}_k,P_k\rangle$ for $1 \le k \le n$. Then for each k, we suppose that the event E_k depends only on the factor Ω_k (this is how one models independence). This means that E_k can be written as

$$E_k = \{\langle\omega_1,\dots,\omega_n\rangle : \forall i(\omega_i \in \Omega_i) \,\&\, \omega_k \in U_k\},$$

for some subset U_k of Ω_k.

Let N be a "nonmeasurable random variable" representing the number of events E_k that happen, i.e., N is a function from Ω to the natural numbers such that $N(\omega)$ equals the cardinality of the set $\{k : \omega \in E_k\}$.

Then we have:

Theorem 2. *Suppose $E_1,\dots,E_n$ are saturated nonmeasurable. Let J be any non-empty proper subset of $\{0,\dots,n\}$. Then the event $N_J = \{\omega \in \Omega : N(\omega) \in J\}$ is saturated nonmeasurable.*

Hence nothing non-trivial can be probabilistically said about the number of the E_k that occur.

The proof of Theorem 2 will use the terminology in Pruss (2013). In particular, for any function f, measurable or not, on a probability space there is a *maximal nonmeasurable minorant* f_* and a *minimal nonmeasurable majorant* f^* such that $f_* \le f \le f^*$ and such that for any measurable functions g and h such that $g \le f$ and $f \le h$ we have $g \le f_*$ almost everywhere (i.e., outside of some set—perhaps empty—of probability zero) and $h \ge f^*$ almost everywhere.

Proof. Define the function Ψ_k on Ω_k by $\Psi_k(\omega) = 1$ if $\omega \in U_k$ and $\Psi_k(\omega) = 0$ otherwise. Define the function Y_k on Ω by $Y_k(\langle\omega_1,\dots,\omega_n\rangle) = \Psi_k(\omega_k)$. Thus, $Y_k(\omega)$ is 1 if $\omega \in E_k$ and 0 otherwise.

Because E_k is saturated nonmeasurable, $(Y_k)_* = 0$ (almost everywhere—I will omit such qualifications from now on) and $(Y_k)^* = 1$. It follows from Pruss (2013, Proposition 1.2) that $(\Psi_k)_* = 0$ and $(\Psi_k)^* = 1$. It follows that U_k is a saturated nonmeasurable subset of Ω_k. Let $\underline{P}_k$ and $\overline{P}_k$ be probability measures on Ω_k extending P to a σ-field that includes E_k and such that $\underline{P}_k(U_k) = 0$ and $\overline{P}_k(U_k) = 1$. These exist by Pruss (2013, Lemma 1.5): just let f in the Lemma be Ψ_k.

Now let j be any member of J. Let Q be the probability measure on Ω formed as the product of the probability measures $\overline{P}_1, \ldots, \overline{P}_j, \underline{P}_{j+1}, \ldots, \underline{P}_n$. Then $Q(E_k) = 1$ if $k \le j$ and otherwise $Q(E_k) = 0$. Thus with Q-probability 1, exactly j of the E_k occur, and since $j \in J$, we have $Q(N_J) = 1$.

Let ℓ be any member of $\{0, \ldots, n\} - J$. Let R be the product of the probability measures $\overline{P}_1, \ldots, \overline{P}_j, \underline{P}_{j+1}, \ldots, \underline{P}_n$. Then $R(E_k) = 1$ if $k \le \ell$ and otherwise $R(E_k) = 1$. Hence, with R-probability one exactly ℓ of the E_k occur, and since $\ell \notin J$, we have $R(N_J) = 0$.

Thus there is an extension of P on which N_J has measure zero and an extension of P on which it has measure one, from which it follows (Pruss 2015, Lemma 2.3) that N_J is saturated nonmeasurable with respect to P. □

References

[1] Edwards, W., Lindman, H. and Savage, L.J. (1963). "Bayesian statistical inference for psychological research", *Psychological Review* **70**:193–242.

[2] Grünbaum, Adolf. (1993). "Narlikar's 'creation' of the Big Bang universe was a mere origination", *Philosophy of Science* **60**:638–646.

[3] Halperin, Israel. (1951). "Non-measurable sets and the equation $f(x+y) = f(x) + f(y)$. *Proceedings of the American Mathematical Society* **2**:221–224.

[4] Hume, David. (1779). *Dialogues Concerning Natural Religion.*

[5] Jeffrey, Ricahrd C. (1969). "Statistical explanation vs. statistical inference." In: Nicholas Rescher (ed.), *Essays in Honor of Carl G. Hempel*, Dordrecht: Reidel, 104–113.

[6] Kleinschmidt, Shieva. (2013). "Reasoning without the Principle of Sufficient Reason." In: Tyron Goldschmidt (ed.), *The Puzzle of Existence: Why Is There Something Rather Than Nothing?* New York: Routledge, 64–79.

[7] Koons, Robert. (2006). "Oppy and Pruss on Cosmological Arguments." Presented at the Central Division Meeting of the American Philosophical Association, Chicago, April 2006.

[8] Leslie, John. (1979). *Value and Existence*. Totowa, NJ: Rowman and Littlefield.

[9] Pruss, Alexander R. (2001). "The cardinality objection to David Lewis's modal realism", *Philosophical Studies* **104**:167–176.

[10] Pruss, Alexander R. (2006). *The Principle of Sufficient Reason: A Reassessment.* New York: Cambridge University Press.

[11] Pruss, Alexander R. (2013). "On the Law of Large Numbers for nonmeasurable identically distributed random variables", *Bulletin of the Polish Academy of Sciences - Mathematics* **61**:161–168.

[12] Rescher, Nicholas N. (1995). *Satisfying Reason: Studies in the Theory of Knowledge.* Dordrecht/Boston/London: Kluwer.

[13] Rescher, Nicholas N. (2010). *Axiogenesis: An Essay in Metaphysical Optimalism.* Lanham, MD: Lexington.

[14] Salmon, Wesley C. (1989). *Four Decades of Scientific Explanation*. Minneapolis: University of Minneapolis Press.

[15] Sierpiński, Wacław. (1938). "Sur un problème concernant les famille d'ensembles parfaits", *Fundamenta Mathematicae* **31**:1–3.

[16] Taylor, Richard. (1974). *Metaphysics*. Englewood Cliffs, NJ: Prentice-Hall.

[17] van Inwagen, Peter. (1983). *An Essay on Free Will*. Oxford: Oxford University Press.

[18] Whitcomb, Kathleen M. (2005). "Quasi-Bayesian analysis using imprecise probability assessements and the Generalized Bayes' Rule", *Theory and Decision* **58**:209–238.