

Skepticism and the Principle of Sufficient Reason

January 8, 2020

Robert C. Koons and Alexander R. Pruss

1. Introduction

Something as fundamental as the Principle of Sufficient Reason cannot be defended by *deducing* it from some still more fundamental rule. It must be defended *dialectically*: by showing that disastrous consequences flow from rejecting it. That is our strategy: to show that denying the *a priori* knowability of the Principle of Sufficient Reason leads to radical empirical skepticism, in the form of denying that we can know that we have any empirical knowledge of the external world (either in particular or in general).

In Section 2, we articulate a form of the PSR that is restricted to basic natural facts. We will argue that all of our empirical data and all of our intermediate empirical inferences fall within this class of natural facts. The PSR we define can be used to demonstrate the existence of a supernatural fact.

We propose in Section 3 our main thesis: that the denial of the *a priori* knowability of this version of the PSR leads to the impossibility of knowing that we have any empirical knowledge. To thwart this kind of empirical skepticism, it is necessary to exclude all skeptical hypotheses from the class of epistemically “close” scenarios. We argue that, in the absence of the PSR, scenarios in which our empirical data occurs without any cause cannot be counted as “distant”.

In Section 4 we lay out reasons for thinking that a principle of causal explanation is needed and that this principle must (if skepticism is avoided) apply to all of the basic natural facts, yielding the form of the PSR introduced in Section 1. We then turn, in Section 5, to six accounts, popular in contemporary literature, of the necessary conditions for empirical knowledge, arguing that in each case the a priori knowability of our version of the PSR is needed. In Section 6, we focus more specifically on the case of probabilistic knowledge.

We then turn to two sets of objections: one based on Neo-Humeanism (Section 7), and one based on the essentiality of origins (Section 8). We also offer a separate and more direct argument for our main thesis in Section 9.

2. Restricting the PSR to *Natural Facts*

One very strong version of the Principle of Sufficient Reason implies that every truth or fact has an explanation. We can go further and assert that all *pluralities* of actual facts have explanations (that is joint or common explanations).

Principle of Sufficient Reason (Unrestricted). If the xx's are some actual facts, then there is an actual fact y that explains the xx's.

It is impossible to assert an unrestricted PSR while requiring non-circular explanations, since the totality FF of all facts must have one of its members as its explanation (since all explanations of

actual facts are actual). Here, we are assuming a principle of comprehension for plurals.¹ Let's also suppose that we can rule out extended circles of explanation (cases of A_1, A_2, \dots, A_n in which A_1 explains A_2 , A_i explains A_{i+1} , and A_n explains A_1 , where $n > 1$) and allow only self-explanatory and partially self-explanatory facts. A fact is self-explanatory if it explains itself, and a fact A is partially self-explanatory iff it is a proper part of a fact that explains A .

The Unrestricted PSR, together with the No Extended Circles principle and the Transitivity of Explanation, entails that some facts are wholly self-explanatory.

This may seem to be a *reductio ad absurdum* of the three principles, since it seems impossible for a fact to be self-explanatory. After all, if x explains y , then y does not explain x , and so if we let $y = x$, we should conclude that we cannot have x explaining x .² However, there are plausible cases of self-explanatory facts: e.g., that $0 = 0$ and that all triangles have three vertices. We can think of a self-explanatory fact as one such that knowing that the fact obtains and fully understanding the nature of the fact removes all of the mystery about why the fact obtains. While our two examples above were analytic, there may even be synthetic self-explanatory facts, such as the fact that water is H_2O .³ Once we understand what it means to say that water is *identical* with H_2O , no mystery about water being H_2O remains.

¹ For reservations on this point, see Pruss and Rasmussen, 2015. For present purposes, we need only the comprehension principle that if some wholly particular, natural facts with feature F exists, then there the plurality of such wholly particular, natural facts exists (see below in this section for the relevant definitions).

² We are grateful to Graham Oppy for this objection.

³ We are grateful to Yunus Prasetya for alerting us to the interesting question of what explains the identity of water with H_2O .

Alternately, we could say that there are two kinds of facts: There are facts that, even when their nature is fully understood, still have a mystery about them and there are facts that when their nature is fully understood are wholly unmysterious.⁴ If explanation is removal of mystery, then facts of the second sort cannot be explained, but only because there is no need to explain them. We could then reformulate our principles and arguments by replacing statements like “x explains y” with disjunctions like “x explains y, or $x=y$ and x is wholly unmysterious.”

Next, observe that there are at least two senses of partial explanation. We can say that x partially explains y provided that x is a part of an explanation of y, or we can say that x partially explains y provided that x explains a part of y. The first sense of partial explanation requires that there be an explanation; the second does not. We will use the language of partial explanation in the first sense. Given this, we have:

No Extended Circles of Explanation. If x partially explains the yy’s, and y is one of the yy’s and not identical to x or any part of x, then y does not partially explain x.

Transitivity of Explanation (Complete and Partial). If x explains the yy’s, and one of the y’s explains the zz’s, then x explains the zz’s, and if x partially explains the yy’s and one of the y’s explains the zz’s, then x partially explains the zz’s.

Theorem 1.⁵ Given the unrestricted PSR, some facts are wholly self-explanatory.

⁴ Cf. “autonomous facts” in Dasgupta 2014.

⁵ Our theorems and lemmas are conditional on previously displayed principles or axioms. We won’t always explicitly specify which ones we are making use of in the statement of the result, so as to avoid undue prolixity, though sometimes we will mention the most salient of these dependencies.

Proof. Consider the plurality of all not wholly self-explanatory facts. Call this plurality PP.

Claim: PP must have at least a partial explanation that is wholly self-explanatory.

Proof of Claim: Suppose for contradiction PP has no wholly self-explanatory partial explanation.

Given the Unrestricted PSR, PP must have an explanation Q, and, by hypothesis, Q is not wholly self-explanatory. Since Q is not wholly self-explanatory, Q is a member of PP. Moreover, since Q is not wholly self-explanatory, Q is partly explained by some R that is distinct from Q. Now, either R is wholly self-explanatory or not. (i) If R is wholly self-explanatory, then R partially explains Q and Q wholly explains PP, and so R partially explains PP. In this case, the Claim follows immediately from the Transitivity of Explanation. (ii) Assume R is not wholly self-explanatory. Then R is a member of PP. But this means that Q wholly explains R and R partly explains Q, in contradiction to No Extended Circles.

So, PP has a partial explanation that is wholly self-explanatory (the Claim). Therefore, there is at least one self-explanatory fact (Theorem 1).

We can consistently assert that every plurality of facts of a certain kind (let's call them *ordinary facts*) has a non-circular explanation.

Restricted (and Non-Circular) Principle of Sufficient Reason. If the xx's are some ordinary actual facts, then there is an actual fact y, not one of nor part of one of the xx's, that explains the xx's.

If the Restricted and Non-Circular PSR (RPSR) is true, then there must exist some extraordinary facts, since the totality of all ordinary facts would have to have an explanation, and that explanation could not be ordinary. This gives us a path to the result in Theorem 0 on other premises.

If we define the ordinary facts as the facts that are not wholly self-explanatory, we can derive the Restricted and Non-Circular PSR from the Unrestricted PSR, given the distributivity of explanation.

Distributivity of Explanation. If y explains the xx 's, and z is one of or part of one of the xx 's, then y explains z .

Claim: Given Distributivity of Explanation, the Unrestricted PSR entails the Restricted (and Non-Circular) PSR.

Proof. Suppose that the xx 's are ordinary facts. By the Unrestricted PSR, there is a fact y that explains the xx 's. Suppose for contradiction that y is one of or part of one of the xx 's. Then, by a principle of Distributivity of Explanation, y wholly explains itself. But then the xx 's are not all ordinary facts. Contradiction.

Let's say that a *natural* being is a being that is wholly composed of things that have intrinsic attributes that are finitely measurable (with some quantitative measure other than zero or

infinity) or that are bounded in some way (including spatial or temporal bounds). A *supernatural* being is one that has at least one simple (atomic) part that is unbounded and infinite in measure. By being ‘part of a fact’, we mean that the existence of such an entity is entailed by the fact, in some relevant sense of entailment (by strong-Kleene three-valued logic or Belnap-Dunn relevance logic, for example). A supernatural fact is one that includes the existence of such a supernatural being as a part. A basic natural fact is any wholly particular fact that is not supernatural.

We claim that no wholly self-explanatory fact is a basic natural fact. Equivalently, we claim that any explanation of any plurality of basic natural facts must be distinct from any member of that plurality. (Proof of the equivalence. Left-to-right, by contraposition. Suppose that some explanation of a plurality of basic natural facts belongs to that plurality. Then that explaining fact is self-explanatory, by the Distributivity of Explanation. Right-to-left, also by contraposition. Suppose that there is some wholly self-explanatory basic natural fact. Then the degenerate plurality consisting of just this fact is an explanation of the plurality, an explanation that belongs to the plurality.) Our claim seems intuitively obvious, since basic natural facts are atomic facts that include finite measures and boundaries that must somehow be explained and do not in any intuitively clear way explain themselves. However, we will not rest with such a naked appeal to intuition. We will argue that such a principle must be accepted on epistemological grounds, on pain of embracing global empirical skepticism.

To summarize, we endorse the following principles and definitions:

Definition of supernatural facts. A fact is *supernatural* iff it entails (in an appropriate relevance logic) the existence of a simple being that is unbounded and wholly infinite in intrinsic measure.

Definition of basic natural facts. A fact is *basic* and *natural* iff it is wholly particular (does not contain any universal generalizations) and not supernatural.

Non-Circular PSR for Basic Natural Facts. If the *xx*'s are basic natural facts, then there is a *y* that is not one of the *xx*'s or part of any of the *xx*'s, and *y* explains the *xx*'s.

If we assume that all explanations of particular facts must include a particular fact, we can use this version of the PSR to demonstrate the existence of at least one supernatural fact, and so the existence of at least one concrete supernatural being.

In the next section, we take up the epistemological argument for the Non-Circular PSR for Basic Natural Facts.

3. The Prima Facie Threat of Skepticism

In order to know that we have any empirical knowledge, we must somehow be in a position to know that the Non-Circular PSR for Basic Natural Fact is true and, not only true, but robustly true, true in all “nearby” or normal possible worlds. First, we must introduce a distinction between epistemically possible scenarios and metaphysically possible worlds. Each *scenario* is a

representation or model of the whole of modal reality. Thus, each scenario includes an actual world and many merely possible worlds (that is, it includes representations of what the actual world and merely possible worlds might be). We are working with a two-dimensional model: a “horizontal” array of epistemically possible scenarios, and, within each scenario, a “vertical” array of supposedly metaphysically possible worlds.

Very plausibly, the knowledge we have in the actual world depends on what happens in nearby possible worlds. Therefore, our knowledge of our knowledge depends on what happens in the possible worlds within each of the epistemically possible scenarios. If we are to know that we have some empirical knowledge, then we must have true empirical beliefs in each epistemically possible scenario that are appropriately related to relevant counterfactual conditions.

We will claim that the modally robust truth of the Non-Circular PSR for BNFs is a presupposition of all empirical knowledge of the external world, and we further claim that our knowledge of this robust truth cannot (on pain of circularity) be justified empirically. Therefore, the robust truth of this PSR must be *a priori* knowable, if we are to know that we have any empirical knowledge of the external world at all.

Central Claim. If we can know that we have any empirical knowledge of the external world (either in general or with respect to some particular proposition), then we are in a position to know *a priori* that the PSR for basic natural facts is robustly true.

Why is the PSR in this form a presupposition of empirical knowledge? If this PSR is false, either in the actual world or in “nearby” worlds in any epistemically possible scenario, then we must take seriously the possibility that some natural facts are uncaused. By considerations of symmetry (which we will discuss in Section 3), if *any* natural fact might lack a cause, then any relevantly similar natural fact might lack a cause. We will argue that any natural fact is relevantly similar to any other. So, if there is some natural fact that might (in nearby worlds) lack a cause, *any* natural fact might lack a cause.

Let’s call any of the sensory or mnemonic appearances, experiences, or impressions upon which we base our empirical beliefs our *empirical data*. If the robust truth of the PSR were not *a priori* knowable, any or all of our empirical data might (for all we can know) be uncaused, that is, are uncaused in an epistemically possible scenario. Obviously, if data are uncaused, they cannot be caused in the right way, in the way required for empirical knowledge of extra-mental facts. This means that there are epistemically possible scenarios in which we lack all empirical knowledge of extra-mental facts but nonetheless have the same sensory inputs. Hence, we cannot know that we have *any* empirical knowledge, beyond our knowledge of the immediate mental data themselves. But we do know this—hence the robust truth of the PSR must be *a priori* knowable by to us.

This is true even if we assume naïve perceptual realism. Suppose that we *know* (in the actual world in the true scenario) that the things we perceive exist and are as they seem. This would translate into the fact that the perceived things exist and are as they seem in the actual world of every epistemically possible scenario. However, this guarantees only that our empirical beliefs

are *true* in those worlds—not that they constitute knowledge. In the absence of the knowability of the PSR, the empirical beliefs will be uncaused in the actual world or in nearby worlds in some of the epistemically possible scenarios.

What if we suppose that our empirical data is partly individuated by its causal connections to the environment?⁶ If this is true, it entails that *if* our data is caused in some way in the actual world, then that same data must be caused in other worlds. However, if we do not know the restricted PSR to be true, there will be *epistemically* possible scenarios in which the data is uncaused, both in the actual world and in modal alternatives. In addition, when evaluating the modal conditions of knowledge in each scenario, we can make use of non-identical counterparts (in David Lewis's sense)—see section 8 below.

But what if the content of our sensory data is determined by its causal connections to the environment? And what if we can know directly (perhaps by introspection) what that content is? Won't this introspective knowledge suffice to ensure that our data is caused in all epistemically possible scenarios? Three points in response. First, all we can know by introspection is that our data has whatever causal background suffices to give it its actual content, and it might be that that background is empty in every case. For all we can know a priori without the PSR, this is compatible without sense data being uncaused. For our introspective knowledge of content is quite limited. Second, this cannot guarantee that all of the links between our sensations and the external facts are caused, and the appearance of any of those links without cause would undermine our empirical knowledge. Finally, this objection, if sound, would prove too much,

⁶ We thank Fatema Amijee for this question.

since it would ensure that all of our empirical beliefs are infallible, and that they must constitute always knowledge. This would contradict the obvious possibility of error and mere true opinion.

In fact, if the PSR cannot be known to be true, there will be epistemically relevant scenarios in which we are in a condition that is at least as bad as the typical skeptical scenarios, like Descartes's victim of the evil demon or the brain in a vat. In the typical skeptical scenarios, our empirical data are caused in the wrong way, by the deceiving demon or the brain manipulating program. Even if these beliefs are always coincidentally true, they cannot constitute knowledge. If the PSR is false in epistemically possible scenarios, then there are relevant scenarios in which our condition is just as bad as they are for Descartes's hypothetical victim—in which we have empirical data, but the data is not caused in the right way to provide empirical knowledge, because the data is not caused at all but has simply “popped” into existence without any cause at all.

Even if the empirical *data* are caused in an accessible scenario, we would lack knowledge in that scenario if any of the causal links *between* that data and its object were to exist uncaused. For example, suppose that the photons by which you see your computer screen were simply to pop into existence without cause, instead of being reflected on or generated by the screen itself. In that case, you would not really *see* the screen at all and so would have no empirical knowledge.

Furthermore, if the PSR fails in epistemically possible scenarios, there will be many cases in which one's empirical beliefs will not only fail to be caused in the right way but will also be false. For example, consider any of our empirical beliefs about the future. If the PSR fails in

possible scenarios, then our actual beliefs about the future could all be falsified in those scenarios, thanks to the intervention of uncaused preventers. We believe that the sun will rise tomorrow, but there would be a possible scenario in which the sun is destroyed today by an uncaused black hole of sufficient size and proximity.

4. Symmetry Considerations and *A Priori* Knowledge

The skeptical scenarios in Section 3 depended on a kind of generalization or epistemic symmetry principle: if there is some natural fact that could (in an epistemically possible scenario) exist uncaused, then *any* natural fact could (in such a scenario) exist uncaused. Why should this be so? Why can't the uncaused events that are epistemically accessible belong entirely to epistemically irrelevant domains, so that no case of empirical knowledge is undermined? Why can't we know that our empirical data and the actual causal links between that data and their objects have causes in all epistemically possible scenarios, without knowing anything as strong and general as the Non-Circular PSR for BNFs?

We've argued that the PSR must be available as *a priori* knowable. So, first of all, we can rule out *de re* knowledge with respect to each empirical datum and each causal link. It is impossible to have *a priori* knowledge about the cause of a particular fact that is not supported by some *de dicto* knowledge, that is, by knowing that all facts of some kind K are caused and that this fact is of kind K. All of our *de re* knowledge, our knowledge by direct acquaintance, is empirical and not *a priori*. Even if we suppose that some of our *de se* knowledge, such as my knowledge that I

exist, is *a priori*, it is surely impossible to know *a priori* that *I* have a cause but that other similar beings do not.

Thus, in order for us to know that we have any empirical knowledge, some form of Restricted PSR must be robustly true (in relation to all *a priori* accessible scenarios):

Generic K-Restricted PSR. If the *xx*'s belong to class K, then there is a *y* that is not a member or part of a member of the *xx*'s and that explains the *xx*'s.

In fact, we not only must we be able to know *a priori* that the K-Restricted PSR is true, but we must also be able to know *a priori* that each of our empirical data and epistemic causal links belongs to the restricted class K to which the Restricted PSR applies, since not only is the K-Restricted PSR a presupposition of all empirical knowledge, but so also is the *knowable applicability* of the K-Restricted PSR to any particular case of empirical data or epistemic link.

In order for the K-Restricted PSR to do the job of satisfying the presuppositions of empirical knowledge, it must meet four conditions:

- (i) The principle must itself be *a priori* knowable, and so class K must be intuitively connected with *aptness to being caused*;
- (ii) it must be *a priori* knowable that nearly all empirical data and nearly all epistemic links (he events or states that mediate causally between those data and their putative external objects) belong to class K, and not by way of supposing that they belong to some proper sub-class;

(iii) the principle must be counterfactually robust (like a law of nature or a metaphysical necessity), and this also must be *a priori* knowable;

(iv) there is no strictly stronger principle satisfying conditions (i)-(iii).

We've added condition (iv) because we want to identify all of the weakest principles satisfying (i) – (iii). We claim that only the K-Restricted PSR in which the class of *basic natural facts* plays the role of class K meets these four conditions:

PSR for Basic Natural Facts. If all of the xx's are basic natural facts, then there is a y that is not one of the xx's that that explains the xx's.

Condition (i) means that we can start with a finite list of potential candidates and proceed by a process of elimination. Interestingly, we don't need to worry that the right principle is one involving a condition that no one has thought of in this connection—it is implausible that we could be completely ignorant of a principle that is fundamentally knowable *a priori* and upon which our empirical knowledge depends.

We can use the first three conditions to derive something like the non-anthropocentricity of class K. That is, the definition of class K cannot involve any essential reference to human beings or indeed to finite, empirical knowers of any kind. It's unreasonable to think that we could know *a priori* that some anthropocentric condition K is associated with a counterfactually robust causal principle, as condition (iii) requires.

Accepting the *a priori* knowability of the K-Restricted PSR as parameterized to basic natural facts is sufficient to avert this particular threat of global skepticism. We can know *a priori* that all of our empirical data are natural, since they are evidently atomic and do not by themselves (either individually or collectively) entail the existence of any infinite or unbounded entity. We can also know this *a priori* about any body of conceivable scientific data, including all experimental results, observations, traces, remnants, and fossils, and about any of the causal links between these bits of data and their putative objects.

In contrast, here are some failed candidates of epistemologically adequate parameters for the K-Restricted PSR:

- (1) The class of non-first events (events preceded in time by other events).⁷
- (2) The class of facts not including an infinite causal regress.
- (3) The class of facts that are not both cosmic in scale and low in entropy.

Each of three corresponding K-restricted PSRs fails to dispel the threat of global empirical skepticism from uncaused sensations. For all I can know *a priori*, it might be that the states of affairs making up my present state of consciousness are all first events. My belief in the reality of the past cannot constitute knowledge *a priori*, without any tacit appeal to my knowledge of causation. Similarly, the events making up my consciousness over a very short period might (for all I can know *a priori*) contain one or more infinite causal regresses, similar to the cannonball

⁷ A principle recommended by Graham Oppy (in conversation).

regress discussed by Co-Author.⁸ Finally, it might well be (for all I can know *a priori*) that my present state of consciousness is the whole of the cosmos. I also cannot know *a priori* that its entropy is very high. Thus, I cannot know *a priori* on the basis of these three principles that my present state of consciousness has any cause whatsoever, much less that it has the sort of cause needed for me to have any empirical knowledge.

5. The Conditions of Knowledge

Here are six necessary conditions for knowledge from the epistemological literature:

1. **Reliabilism.** A process can generate knowledge only if it has a high and determinate objective probability of producing true beliefs.
2. **No-undercutting-defeater.** A subject S knows that p only if there is no fact q such that (1) that q is a good reason (*prima facie*) for thinking that S's actual grounds/basis for believing that p are not reliably truth-ensuring, and (2) there is no r known to S in a way that is not *prima facie* defeated by q such that that r is good reason for thinking that q is not in fact (all things considered) a defeater for S's knowing that p.⁹
3. **Nozick Sensitivity.** A subject S knows that p only if it is objectively very likely that if p were false, S would not believe that p.
4. **Nozick Converse Sensitivity.** A subject S knows that p only if it is objectively very likely that, were S to believe that p on exactly the grounds on which S does in fact believe that p, p would still be true.
5. **Safety (Williamson).** A subject S knows that p only if in all relevantly close world w in which S believes p on the grounds/basis on which S does in fact believe p, S does not falsely believe that p. (Williamson 2000)

⁸ Pruss 2006, 44-6.

⁹ See Koons 2018, 251-3.

6. **High Objective Chance (Hawthorne-Lasonen-Aarnio).** If the objective chance of the truth of p is not determinately high at t , then no subject can know that p at t . (Hawthorne and Lasonen-Aarnio 2009)

In each case, the symmetrical falsity of the Restricted PSR in all epistemically accessible scenarios is enough to falsify the condition, and hence acceptance of any one of these requires the convinced anti-sceptic to accept the Restricted PSR for BNFs. We will take up the case of High Objective Chance in the following section (Section 6).

5.1 Reliabilism

If, in some epistemically possible scenario, the PSR for BNFs is false in nearby possible worlds, then the reliability of empirical methods of belief formation has an indeterminate probability in that scenario, since the intermediate steps of the process could occur uncaused, and there is no way to put any upper bound on the probability of such uncaused events. In fact, it is impossible to assign any objective probability to classes of uncaused or unexplained events.

An objective probability is either a causal propensity or some kind of ideal relative frequency. In both cases, in order to determine the objective probability of an event of a certain kind, it is necessary to take as background condition some set of events that are causally and explanatorily prior. It is only relative to such an explanatorily prior background that it makes any sense to talk of the objective probability of an event. Hence, uncaused events cannot have such probabilities.

An *objective* probability is either a propensity for a result (in each single case) or some kind of actual or hypothetical relative frequency. You can't, by definition, have a propensity to produce an explained event. So, the objective probability of unexplained events would have to be some kind of ideal relative frequency. But when we think of objective probabilities, the relative frequency we have in mind is the frequency of a certain result *given* a causally prior condition of a certain kind. And without a causally prior condition, this seems to make no sense and yields an undefined probability.

In response, one might ask: why couldn't we discover the objective probability of uncaused events by comparing the number of uncaused events with the total number of actual events, both caused and uncaused? Suppose in a possible world w , there are m events that occur in the actual world, n of which are uncaused. Couldn't we take n/m to be the objective probability of the occurrence of an uncaused event in such a scenario?

In order for such an actual frequency to give us an estimate of an objective probability, we would have to suppose that, in a given epistemically accessible scenario, all worlds that are close to the actual world have frequencies that are very close to n/m . But what could possibly explain why all of those nearby worlds would instantiate ratios that approximate n/m ? Since the occurrence of unexplained events is by hypothesis unexplained, there can be no explanation for why there are not more or fewer such events, and so no basis for excluding worlds in which the ratio approaches 1 from the neighborhood of w .

In addition to this fundamental problem, there are also some important technical difficulties with this proposal. First, can we really speak of the total number of actual events in the actual world and nearby worlds? We cannot, if space and time are even potentially continuous, or if space or time are potentially infinite in extent. So, the opponent of the PSR would have to suppose that actual space and time, and space and time in all nearby worlds, are finite and have uniform bound in size, both in composition and extent.

Let's suppose, for the sake of argument, that there are only a finite number of spatiotemporal slots in which possible events can be located. Even so, we are not home free, since we have to consider all of the actually occurring absences (negative events and conditions). If the number of properties and kinds of things is infinite, then there are infinitely many actual events even in a spatiotemporally finite world. To make the proposal work, we would have to assume that there are, as a matter of metaphysical necessity, only a finite number of families of properties, each of which is composed entirely of mutually exclusive kinds. On this hypothesis, there would be only finitely many negative facts, since the instantiation of a member of one of these families in a spacetime slot would be the truthmaker for the non-instantiation of all of the other members of that family in that slot.

Even under these very favorable assumptions, we would face one last obstacle. In order for our empirical methods to be reliable, it is not enough for there to be some upper bound on the probability of uncaused events in general. What we have to consider is the probability that an event of forming a certain belief (event of type K) occurred uncaused. Obviously, the frequency of the occurrence of an event of type K is quite low: the number of events that involve the fixing

of a human empirical belief is small compared to the total number of events. If the probability of the occurrence of a caused instance of K is ϵ , and the probability of the occurrence of an uncaused event is δ , then the conditional probability of an event of type K's being uncaused will be approximately $\delta/(\epsilon + \delta)$, by an application of Bayes's Theorem:

$$\begin{aligned}
 & \text{Prob}\left(\frac{E \text{ is uncaused}}{E \text{ is of type K}}\right) = \\
 & \frac{\text{Prob}\left(\frac{E \text{ is of type K}}{E \text{ is uncaused}} \cdot \text{Prob}(E \text{ is uncaused})\right)}{\text{Prob}\left(\frac{E \text{ is of type K}}{E \text{ is uncaused}} \cdot \text{Prob}(E \text{ is uncaused}) + \text{Prob}\left(\frac{E \text{ is of type K}}{E \text{ is caused}}\right) \cdot \text{Prob}(E \text{ is caused})\right)} \\
 & \approx \frac{\delta}{(\delta + \epsilon)} \approx \frac{\delta}{(\delta + \epsilon)}
 \end{aligned}$$

We assign the value 1 to $\text{Prob}(E \text{ is of type K}/E \text{ is uncaused})$, because there is no determinate value less than 1 that such a probability could take. The ratio $\delta/(\epsilon + \delta)$ could be quite large, given the smallness of ϵ , far too large for our belief-forming capacities to count as reliable. In order to avoid having a defeater to any given empirical belief, we would have to know a priori that delta is extremely low (see 5.2).

Could there be an objective probability, not that some unexplained event or other occur, but that an unexplained event of a particular *type* occur? If so, there would be different probabilities for different types, in certain cases. What could possibly ground these differences? If an observed frequency of unexplained events of a certain type is to constitute a counterfactually robust probability, there must be some explanation for why the actual frequency is what it is, an explanation robust enough to ground similar frequencies of the same type of event in nearby

worlds. Each type of event would be tied to a different explanation. If such were the case, the occurrence of the events of a given type would not be unexplained—they would have a probabilistic or stochastic explanation.

5.2 No Undercutting Defeater

If, in some epistemically possible scenario, the PSR is false in some nearby world, then this very fact constitutes in that scenario a defeater for S's knowing anything empirically, since it provides good *prima facie* grounds for doubting that the methods on which S is relying are reliable.

In addition, there can be in such a scenario no defeater-defeaters of this defeater. That is, there are no grounds for defeating this defeater that are not themselves defeated by the possible falsity of the PSR. If the PSR is false in nearby worlds, this fact is a *prima facie* defeater of any bit of knowledge sufficient to defeat the defeater. Any knowledge we have of the actual causal antecedents of any of our empirical beliefs must itself be empirical, and so subject to the same *prima facie* defeater. And empirical knowledge is impossible in the presence of an ultimately undefeated defeater.

Consequently, if the PSR fails to be robustly true in every epistemically scenario, we cannot know that we have any empirical knowledge.

To make this more concrete, consider the grounds for doubting the reliability of our empirical data (as considered in section 5.1). Suppose, for the sake of argument, that our empirical beliefs

are in fact caused in a reliable way. Assume further that we cannot know *a priori* that the PSR is robustly true. Hence, there will be epistemically accessible scenarios (relative to our possible *a priori* knowledge) in which there are possible worlds close to the actual world in which a large proportion of our empirical data occurs uncaused. The existence and proximity of such causally unruly worlds constitute a defeater for our empirical knowledge in these epistemically possible scenarios, since they render our empirical data unreliable (with respect to truth). This defeater cannot be defeated in any of those scenarios, since *all* empirical data is *prima facie* defeated by the unreliability of our data. Hence, we would have no empirical knowledge in any such scenario, and, consequently, we cannot actually know that we have any empirical knowledge.

5.3 Nozick Sensitivity (Nozick 1981)

In the absence of the robust truth of the PSR, if p were false, there is an indeterminate probability (see Section 5.1) that S would still believe p , as the result of an uncaused event resulting in S 's believing that p (on the same grounds/basis as S did). Thus, if there are epistemically accessible scenarios in which the PSR fails in any $\sim p$ world, we lack knowledge of p in those scenarios.

If we cannot know the PSR *a priori*, for every empirically known p , there will be such a knowledge-falsifying scenario. Hence, we could not know that we have any empirical knowledge at all.

5.4 Converse Nozick Sensitivity

Again, if the PSR were false in nearby worlds, S might have believed that p on exactly the same grounds as an uncaused event and in the absence of the fact that p. Thus, if there are epistemically accessible scenarios in which the PSR fails in some non-actual worlds in which S believes that p, we lack knowledge of p in those scenarios.

So, once again, the absence of a priori knowledge of the PSR deprives us of the knowledge that we have any empirical knowledge.

5.5 Safety (Williamson)

As we have argued, the nearby possibility of uncaused events brings worlds in which S comes to believe p on the same grounds/basis did, despite the falsity of p, close to the actual world. This violates Safety in those accessible scenarios in which the PSR fails to be robustly true. This will be true for any empirically based belief, and so we will be unable to know that we have any empirical knowledge.

Objection: if the process leading to S's belief that p involved no uncaused events, then any world containing uncaused events is *ipso facto* not close. According to Williamson, our intuitions about closeness should be guided by our intuitions about knowledge, since the concept of knowledge is unanalyzable.

Reply: such a stipulation about closeness would seem arbitrary and unmotivated. If uncaused events can occur at any point with indeterminate probability, then we must always treat them as

real, salient possibilities. A Williamson-like stipulation would involve putting our heads in the sand.

6. Probabilistic Knowledge

In the absence of the PSR, we could not even say that uncaused events are improbable. An event is improbable only if its potential causes are such as to produce the event in question only in exceptional cases. To assign an objective probability to an event is to ascribe a certain kind of cause to it, and, as a result of de Finetti's theorem (as explained by Skyrms 1984, 37-62), it is impossible to assign subjective probabilities coherently without a tacit commitment to objective probabilities.

One of us has recently argued that, again in the absence the presumption of a PSR (Pruss 2017), we cannot confirm or disconfirm any probabilistic hypothesis on the basis of statistical information. This would undermine all of modern science, including our most fundamental physical theory, quantum mechanics.

If a particular sequence of data could occur uncaused and unexplained, there are three possibilities to consider: (I) there is some definite probability that such a sequence would occur on the hypothesis that the sequence is uncaused, (II) there is no definite probability but merely a range of probabilities for such an occurrence, or (III) there is absolutely no probability that can be assigned to such an occurrence. If we suppose case (I), then there is some probabilistic causal explanation that cannot be distinguished empirically from the no-explanation case, and we would

therefore never be able to raise the probability of the probabilistic hypothesis above $1 - p$, where p is the probability of the no-cause hypothesis. Nor is there any *a posteriori* data from outside the sequence that would seem to be relevant to ruling out the no-explanation hypothesis, because there is no reason to ascribe any kind of correlations between a no-explanation hypothesis and other data. The situation is even worse in cases (II) and (III). Therefore, if statistical inference is capable of approaching certainty as the size of the data approaches infinity we must be able to exclude the no-explanation hypothesis on *a priori* grounds.

Could we, however, perhaps suppose that while statistical inference would not approach certainty, it would get close enough for practical purposes, say 0.99999999? Even this seems unavailable--and hence case (I) (and, for similar but more involved reasons, (II)) is not in the end tenable. This is because there are no probabilities in the absence of chances and no chances in the absence of explanations or causes.

This claim about the lack of probabilities conflicts, however, with intuitions related to the principle of indifference. We explain the point through an example. Consider ten coins that come into existence with no explanation, and ask how likely it is that they would all appear heads-up, given that they all popped into existence causelessly. There is a strong intuition that the chance of this is something like $1/2^{10}=1/1024$ (and because this is so low, we might then prefer a hypothesis on which there is a law of nature that coins are always heads up on the basis of the data), and hence the probability that ten coins would pop into existence *all heads up* is less than $1/1024$ (since this probability equals $1/1024$ times the probability that ten coins would

causelessly pop into existence). After all, there is no reason to favor any particular arrangement of heads and tails among the coins that popped into existence, and there are 2^{10} arrangements.

But this reasoning is fallacious. If we assume that there *is* a probability to be assigned to the arrangement *and* that no particular arrangement is favored over any other, then we can say that all arrangements must have equal probability and the probabilities corresponding to the 1024 heads/tails arrangements must add up to 1, and hence each probability must be $1/1024$. But this argument only works if we assume from the outset that there *is* a probability to be assigned to the arrangement. If there is no probability, then we can consistently say that the 1024 possible arrangements do not differ with respect to probability: they all equally have no meaningful probability.

Moreover, more simply, but in greater detail, let us consider a single coin that pops into existence with no explanation at all in a gravitational field. That coin can pop into existence in one of three orientations: heads-up, tails-up, and on-the-side. We can characterize these orientations as follows. We consider a vector coming out of the heads side of the coin at right angle to the coin surface. When that vector is at 90 degrees to the gravitational force, the coin is on its side. When the angle of that vector to the gravitational force is between 0, inclusive, and 90 degrees, exclusive, the coin is tails-up. When the angle is between 90 degrees, exclusive, and 180 degrees, inclusive, the coin is heads-up. For simplicity--but even this can be questioned--let us neglect the case of the coin being on its side. Why think that the angle between the heads-up vector and the gravitational field is equally likely to be between 0 and 90 as between 90 and 180 degrees? Admittedly, there is an equal number of angles between 0 and 90 and between 90 and

180. But there is also an equal number of angles between 0 and 45 and 45 and 180. In both cases, the number of possible angles equals the cardinality of the continuum.¹⁰ It is only because of the nature of the geometry of space that we want to say that the range between 45 and 180 is more expansive than that between 0 and 45: the cardinal number of angles in each range is the same. But the geometry of space is intimately tied to the laws of nature obtaining in space: this is a central insight of Einstein's theory of relativity. And, by definition, coins popping into existence without causal explanation are not governed by the causal laws of nature, and so there is no reason to think that they would have probability distributions tied to the geometry of space, and hence no reason to think they would equally often pop up with orientations to the gravitational field in the 0 to 90 degree range as in the 90 to 180 range, or that they would more often pop up with orientations between 45 and 180 than between 0 and 45. If the probability distributions were tied to the geometry of space, then the regions of space would be part of a causal explanation for the coin-toss results, contrary to hypothesis. So, we cannot even say that the single coin has probability 1/2 of popping into existence heads-up and 1/2 of doing so tails-up when there is no explanation of these events.

7. Objection from Neo-Humeanism

Suppose that we are Neo-Humeans, like David Lewis. Whether we think that possible worlds are concrete universes or some sort of abstract object, we would suppose that the relevant notion of distance between worlds is a matter of human convention and linguistic pragmatics. In particular,

¹⁰ There is, after all, a one-to-one map between the range from 0 to 45 and the range from 45 to 180: a number x between 0 and 45 can be one-to-one paired with the number $180-3x$.

Lewis argued that whether or not the same laws of nature are verified in two worlds is highly relevant to their closeness.

Let's further suppose that the world's actual laws of nature (thought of now as brute regularities that count as laws because of their belonging as axioms to the "best system of laws" for the Humean mosaic of the actual world) are of such a character as to ensure that all of our empirical data and epistemic links have causes. Consequently, these events must have causes in all nearby possible worlds, and the preconditions of empirical knowledge are met.

This objection is sound, as far as it goes. If Neo-Humeanism were the true and adequate account of modal reality, empirical knowledge would not depend on anything like the PSR. This is not surprising, given the extremely deflationary attitude of Humeans toward causation and explanation. There are, however, a large number of compelling objections to the Neo-Humean project, a few of which we will recapitulate here.

The Humean's claim that conformity to the actual laws of nature is crucial in measuring a world's distance from the actual world is not self-evidently correct but is in need of some justification. Since Humeans deny the fundamental reality of causal powers, they ought to measure closeness of worlds solely in terms of how similar their Humean mosaic of qualities are. For this reason, a world in which the pattern of qualities in the Humean mosaic is close to that of the actual world should be relevantly "close" (for the purposes of sensitivity or safety), regardless of whether that world verifies the laws of nature of the actual world. Hence, there will be close worlds in which inductively supported generalizations are falsified in the subject's

immediate future. Consequently, even if it is “rational,” in some sense, to believe empirically confirmed generalizations, we could never know them to be true, if Humeanism were correct.

The laws of nature in our world are in fact grounded in the actual causal powers of existing things, and in the nature of kinds of thing in which those causal powers are grounded: the laws and the natures and powers they are grounded in govern how existing things interact. To detect the laws of nature, we should not look at the axioms of the “best system” of laws: we should instead look to our best account of the natures of things and the powers those natures confer (as Nancy Cartwright has argued). Since the laws of nature govern only how things of various natural kinds interact, they cannot prevent the emergence of new, uncaused things or uncaused changes in existing things. Only a metaphysical principle like the Principle of Sufficient Reason can do that. Laws of nature govern the realm of the caused: they have no purchase on the uncaused.

Consequently, the Neo-Humeans’ assumption that the actual laws hold in all nearby worlds depends implicitly on the robust truth of the PSR. Without such a principle, there is no good reason for supposing that deviations from the laws of nature might not occur in the epistemically relevant neighborhood of the actual world.

The inadequacy of the Neo-Humean account can also be seen by considering the evaluation of the relevant counterfactuals. As one of us has argued (Pruss 2003), the Humeans have no good account of time’s arrow, since the actual laws of nature are time-reversible. The correct evaluation of counterfactuals requires some sort of temporal asymmetry, in which we typically

hold the past constant and let the future vary widely from the actual world. A causal-powers ontology can explain this asymmetry easily. All causal powers have a built-in temporal directionality: an active power is a power to produce new conditions in the future, and a passive power is a capacity to undergo certain changes in response to the exercise of active powers. Moreover, when measuring the closeness of worlds, we much prefer leaving some events miraculously effectless to leaving any events causeless. Otherwise, it would make sense to leave the events in the immediate future of some counterfactual event exactly like the actual world despite the absence of a cause, rather than rendering some actual past events effectless in order to introduce the counterfactual event into the stream of history.

If we adopt a causal-powers ontology as the most viable alternative to Neo-Humeanism, we will need to include the robust truth of the PSR in our theoretical package. In the absence of the robust PSR, causal powers would be toothless, since changes and states could occur causelessly, in complete independence of the array of causal powers in the situation. In addition, the causal powers of particular things could change causelessly, with the result that all causal explanations would be one-off and ad hoc. We could explain any event by positing the required causal powers in the relevant entities, without implying any general truth about explaining similar events in similar circumstances. All causal explanations would be like the appeal to Molière's *dormitivity*.

8. Objection from Origins Essentialism

Suppose we adopt, following Saul Kripke's lead in *Naming and Necessity* (Kripke 1980), a very strong doctrine of origins essentialism, according to which no actual event can exist in any world

in which it is not caused in exactly the way it is caused in the actual world. In addition, let's suppose that kinds of things, including human beings and beliefs, can only exist in worlds in which they are caused to exist exactly as they are caused to exist in the actual world.

Under such assumptions, it would seem to be relatively easy to satisfy the conditions of reliability, sensitivity, and safety, at least with respect to beliefs about the present and past. First of all, assuming some form of functionalism, uncaused mental states would not count as *beliefs* or *sensory impressions* at all, since they would not be caused in the right sort of way.

Consequently, belief-producing processes would be just as reliable as they are in the actual world. Nozick sensitivity and converse sensitivity and Williamson's safety condition couldn't be easily falsified, because uncaused states wouldn't count as *beliefs*.

Our reply would be that, in evaluating possible cases of knowledge, it is sufficient to find states in nearby worlds that count as *counterpart* (in David Lewis's sense) of actual beliefs and sensory or mnemonic impressions, even if those states don't count as beliefs or sensory impressions in those other worlds. Otherwise, any true belief about that belief-state's causal history would automatically count as knowledge. For example, suppose that John truly believes on the basis of unreliable testimony that Bruce and Margaret are his parents. Given origins essentialism, it is impossible for John to believe that proposition in any world in which it is false. This renders Nozick sensitivity and converse sensitivity and Williamson safety all vacuously true, yet clearly the state satisfies only the letter and not the spirit of these three conditions. Here are better versions:

- **Nozick Sensitivity.** A subject S knows that p only if it is objectively very likely that if p were false, no counterpart of S would believe that p (or be in a state that is a counterpart of the state of believing that p).
- **Nozick Converse Sensitivity.** A subject S knows that p only if it is objectively very likely that, were S or some counterpart of S to believe that p (or be in a state that is a counterpart of believing that p) on exactly the grounds on which S does in fact believe that p, p would still be true.
- **Safety (Williamson).** A subject S knows that p only if in all relevantly close world w in which S believes p (or some counterpart of S is in a state that is a counterpart of believing that p) on the grounds/basis on which S does in fact believe p, no counterpart of S is in a state that is a counterpart of falsely believing that p.

For example, suppose that you are in a situation in which your brain might very well have been placed (without your knowledge) in a reality-simulating vat (i.e., such a vat world is relevantly close). And suppose that brains in vats do not have states that count as beliefs or sensory impressions. If you groundlessly believe that your brain is not in a vat, then you do not know this to be so, simply because if it were false, you would not believe it. If it were false and your brain were in a vat, you would be in a state that is a counterpart to falsely believing that you are not a brain in a vat, and that fact should suffice to deprive you of knowledge.

Similarly, if you don't have adequate grounds for thinking that your current states are causeless, you don't know that they have causes, even if it is true that were your current state to occur causelessly, you wouldn't still *believe* that it had a cause. You would be in a state that would be a counterpart of such a false belief, and, if such a causeless state occurs in a relevantly close world, this is enough to deprive you of knowledge in this world.

9. A More Direct Argument

Here is a more direct argument from the denial of skepticism (about the future) to the existence of an explanation for the plurality of natural facts. Suppose that the plurality of natural facts has no explanation. Then that totality exists in the actual world as an unexplained brute fact. If so, then there are relevantly close worlds in which a *slightly* different totality exists instead. In these close worlds, every natural kind has almost the same causal powers as they have in the actual world, or things belong to almost the same natural kinds (as measured by their causal powers) as they belong to in the actual world, or there are counterparts to all the actual things that belong to almost the same natural kinds as the actual things belong to.

In any case, the things and their causal powers are close enough that the history of these worlds up to the present moment is superficially indistinguishable from the actual world, and yet some of these worlds differ significantly from the actual world in the immediate future of any given subject. The deviations are sufficient to deprive actual subjects of any safe beliefs about future events. Hence, total skepticism about the future follows.

If we are evaluating the safety of beliefs about the future, we cannot include the verification of those beliefs in measuring the relevant closeness of worlds, or else we would lose the distinction between knowledge and true belief, or (at least) we would trivialize the safety condition.

Therefore, in evaluating beliefs about the future for closeness, we should consider only facts about the past and present in evaluating closeness. The future-falsifying worlds are *very* close to the actual world, in the absence of a necessary first cause.

So, if we assume that we have some knowledge about the future, then the totality of natural facts must have an explanation.

10. Conclusion

Without some sort of a Principle of Sufficient Reason being true, it is not possible to know that we have any extra-mental empirical knowledge, since there is no way of making the hypothesis that all our mental states are causeless events be even *unlikely*, much less impossible. This yields a transcendental argument for such a PSR.

What if we gave up the claim that we can know that we have any empirical knowledge? Perhaps we should simply treat the hypothesis that we have empirical knowledge as a reasonable conjecture or presumption. Four responses. First, we have shown that we cannot have empirical knowledge of the external world unless *some* form of the restricted PSR is true. If we give up the claim that we can know that we have any empirical knowledge, we will lose our deductive argument (section 4) for the PSR restricted to basic natural facts. However, we can still argue for this particular form of the PSR as the best (i.e., simplest) explanation for our empirical knowledge: for surely we *do* have empirical knowledge, even if we do not *know* that we have it.

Second, even if we suppose that we cannot know that we have (in general) some empirical knowledge or other, on the grounds that considering such a hypothesis invites the consideration of skeptical hypotheses, it remains very plausible that there are particular cases of empirical

knowledge (such as the knowledge that I am now typing on a keyboard) which we can know that we have. However, even such particular instances of second-order knowledge will be impossible unless we can know that PSR for basic natural fact *a priori*, since in each such case we would have to take seriously scenarios in which all our data is uncaused.

Third, whatever epistemic status we give to the hypothesis of our having empirical knowledge, we must give one at least as good also to the PSR. If the former is a reasonable presumption, then so is the latter.

Fourth, if, for each empirical proposition p , we cannot know that we know that p , then that fact is itself a defeater of all of our (first-order) empirical knowledge, since it means that we must always treat skeptical scenarios as epistemically live possibilities. We are not appealing here to the KK principle: we're not assuming that, for *every* p , to know that p , it is necessary that one know that one know that p . As Williamson has argued, one's knowledge that p might be a borderline case of knowledge. However, it is plausible that in order to know that p empirically, there must be some q such that one can know that one knows that q empirically. If scenarios in which I have no empirical knowledge are always epistemically available, this should undermine my first-order empirical knowledge.

We have argued that a PSR restricted to pluralities of ordinary facts is enough to show that there is an extraordinary and therefore *supernatural* fact. Consequently, non-skeptical epistemology requires serious metaphysics, both on the sides of principles and on the side of what exists.

References

Dasgupta, S. (2014). The Possibility of Physicalism. *Journal of Philosophy*, 111(9-10), 557-592.

Hawthorne, J. and Lasonen-Aarnio, M. (2009). Knowledge and Objective Chance. In P. Greenough and D. Pritchard (Eds.), *Williamson on Knowledge* (pp. 92-108). Oxford: Oxford University Press.

Koons, R. (2018). The Argument from Intuition. In T. Dougherty and J. Walls (Eds.), *Two Dozen or So Arguments for God's Existence: The Plantinga Project* (pp. 238-57). Oxford: Oxford University Press.

Kripke, S. (1980). *Naming and Necessity*. Cambridge, MA: Harvard University Press.

Nozick, R. (1981). *Philosophical Explanations*. Cambridge, MA: Harvard University Press.

Pruss, A. (2003). David Lewis's counterfactual arrow of time. *Noûs*, 37(4), 606–637.

Pruss, A. (2006). *The Principle of Sufficient Reason: A Reassessment*. Cambridge, UK: Cambridge University Press, 2006.

Pruss, A. (2017). The Principle of Sufficient Reason and Probability. *Oxford Studies in Metaphysics*, 10, [https://DOI:10.1093/acprof:oso/9780198791973.003.0011](https://doi.org/10.1093/acprof:oso/9780198791973.003.0011)

Pruss, A. and Rasmussen, J. (2015). Problems with plurals. *Oxford Studies in Metaphysics*, 9,
[https://DOI:10.1093/acprof:oso/9780198729242.003.0004](https://doi.org/10.1093/acprof:oso/9780198729242.003.0004)

Skyrms, B. (1984). *Pragmatics and Empiricism*. New Haven, CT: Yale University Press.

Williamson, T. (2000). *Knowledge and its Limits*. Oxford: Oxford University Press.