

Our A Priori Knowledge of the Scope of Causation

Rob Koons

October 5, 2007

1 Introduction

In the early modern period (from Descartes through Kant, at least), the epistemological need for a principle of general causation was a commonplace. However, the collapse of Cartesian foundationalism in the first half of the twentieth century resulted in a general neglect of the role of such a causal principle in accounting for our empirical knowledge. I will not argue for a return to Cartesian assumptions: I accept a broadly common sense, Reidian or Moorean approach to epistemology. However, I will argue that the need to secure knowledge against the possibility of undermining defeaters leads to the restoration of the principle of general causation to an exalted status in epistemology.

Any plausible principle of general causation must fall short of absolute universality. If everything must have a cause, then this would apply to the whole of reality, including the supposed First Cause itself. A coherent principle must

admit of exceptions. However, not all exceptions are equal in respect of their epistemological implications. I will argue that there are strong epistemological constraints that reduce to a narrow band the range of permissible principles of causation. This fact may have great significance for the contemporary debate over the cosmological argument of natural theology: a principle with exceptions narrow enough to escape global skepticism may also be strong enough to serve as the heart of a successful argument for God's existence.

2 Restrictions on the Principle of General Causation: The Need and the Candidates

As I mentioned in the introduction, an axiom asserting the existence of a cause of *every* actual situation is incompatible, in the presence of plausible principles of causation, with the existence of *any actual* situations. If there were any actual situations, there would be a maximal situation, Reality, that was the aggregate of all actual situations. Reality would have to have a cause, which would have to be both actual and wholly separate from Reality, a conclusion that is inconsistent with the definition of Reality.

Therefore, the principle of general causation must be restricted. But how? Here is a list of seven possible candidates that have been proposed:

1. All non-first situations have causes (Graham Oppy).^[7] A situation has *non-firstness* just in case there is an actual situation wholly located at an earlier time than it is.

2. All situations with finite temporal duration have causes.
3. All situations that don't occur at a first moment of time have causes.
4. All situations that don't include infinite causal regresses have causes.
5. All situations that aren't both extremely simple and cosmic in scale have causes.
6. All situations that could (de re) be caused have causes.
7. All situations that are wholly contingent have causes.[4],[5]

The crucial question is this: what are the implications of these and all similar exceptions for empirical epistemology? I will argue that all but the last of these lead inevitably to global skepticism.

3 Two Possible Connections between Causation and Empirical Knowledge

Alvin Plantinga's proper-function epistemology could be summarized by the assertion that there are at least four necessary conditions for knowledge: a state of knowing p must be:

1. a true belief that p ,
2. formed by cognitive processes that are functioning normally, and in circumstances for which they were designed,

3. formed by cognitive processes that have the proper function of producing and sustaining true beliefs, and
4. immune to internal defeaters.

This Plantingan account seems to be fundamentally sound, at least as a good first approximation.^[8] Given this account of knowledge, there are two ways in which a principle of general causation could be connected to the possibility of empirical knowledge: (i) as an objective fact needed as the ground for the reliability of our cognitive processes, and (ii) as a subjectively required presumption needed for immunity to internal defeaters. Let's explore each of these in turn.

Suppose that there were a significant probability that our empirical beliefs occur without cause, or that some of the states that intermediate between our beliefs and their objects so occur. In such a world, the occurrence of those belief states would not reliably indicate their truth, since the occurrence of an uncaused state would be probabilistically independent of the occurrence of those states which might have caused it.¹

I would to introduce here the term 'knowledge-net' for the sake of simpler exposition. A person's knowledge-net consists of all of his² belief-states, to-

¹My argument is intended to be entirely neutral between the various competing metaphysical accounts of the nature of causation, whether nomological-deductive, regularity-based, counterfactual, singularist or what have you. It may be that the conclusion of the argument will have significant implications for the metaphysics of causation (indeed, I am quite certain this is so), but I am being careful not to beg any such questions in the setup of the argument itself.

²For the sake of simplicity, I use the generic masculine pronoun throughout.

gether with those states, if any, that both cause one of those belief-states and intermediate causally between it and its object (or between it and the common cause of it and its object). For example, if S has perceptual knowledge of the fact that p by vision, then S 's knowledge-net includes his belief that p together with those states that are causally intermediate between the fact that p and his belief that p , such as the reflection of light by the objects involved in the fact that p , the transmission of that light to S 's eyes, the occurrence of nerve signals between S 's retina and brain, and S 's visual impressions as of the truth of p .

S 's actual knowledge-net is a part of the actual world that varies from one possible world to another. In order for S to have empirical knowledge, it must be that a principle of general causation applies with high probability to the situations that make up S 's knowledge-net. For this judgment of high probability to be knowledge-conferring, it must not be merely accidental or coincidental that nearly all the situations in S 's knowledge-net are caused. Therefore, it must be at least nomologically necessary (if not metaphysically necessary) that situations of the kinds making up S 's actual knowledge-net be caused, i.e., nomologically impossible for situations of those kinds to occur without a cause.

However, although there must be something like a law of nature guaranteeing that the scope of causation include nearly everything in S 's knowledge-net, and this law of nature must apply to those situations non-accidentally, it could be that the law applies to those situations by virtue of some feature that is essential to S as a particular individual: that is, the applicability of the principle of causation to S 's knowledge-net could well turn on features of that net that are

necessary de re of S himself. Since it is not at all obvious what features might be necessary de re of S , it is difficult to know what exactly the scope of causation must be in order for S to be de re reliable with respect to the formation of his beliefs.

There is, however, a second possible connection to explore: the need for a subjective ground of immunity to internal defeat, in the form of a justified presumption of a principle of general causation. This, I will argue, is a much more promising route to take.

4 Immunity to Defeat as a Necessary Condition for Knowledge

4.1 What are Internal Defeaters?

A *defeater* is a belief that rationally *defeats* one's grounds for believing another proposition. The notion of a defeater was introduced to epistemology by Roderick Chisholm and plays a central role in the epistemologies of John Pollock and Alvin Plantinga.^{[9][8]} In this section I will give my own account of defeaters, drawing heavily on Plantinga's discussion in *Warrant and Proper Function*.

First, following Pollock, we can distinguish two kinds of defeaters: rebutting and undercutting. A rebutting defeater for a belief that p constitutes grounds for believing p to be false; and undercutting defeater does not do so but instead provides grounds for believing that one's disposition to believe that p is epistemically faulty.

Definition 1 A proposition q is a rebutting defeater of proposition p for S iff S believes q and q provides S with adequate grounds for judging that p is false, even when combined with S 's evidence for p .

Definition 2 A proposition q is an undercutting defeater of proposition p for agent S iff S believes q and q provides S with adequate grounds for judging that it is not highly likely that the processes that led to his^{*3} disposition to believe p are warrant conferring.

As Pollock and Plantinga observe, defeaters can themselves be defeated. I will call a defeater defeater a *neutralizing* defeater:

Definition 3 A proposition r is a neutralizing defeater of q in relation to p for S iff S believes r and q , q is a *prima facie* defeater of p for S , and $(r \ \& \ q)$ is not a *prima facie* defeater of p for S .

To ignore the epistemic upshot of one's defeaters is itself a serious cognitive malfunction. Evidently, our *design plan* (insofar as that plan is aimed simply at truth) includes an unvarying disposition to give up beliefs that have been defeated, unless those defeaters have themselves been neutralized by further beliefs. Respect for rebutting defeaters embodies a kind of principle of total evidence: one must not form beliefs based on an arbitrarily selective consideration of available evidence.

Respect for undercutting defeaters embodies a commitment to good probabilistic reasoning and a refusal to rely blindly on a supposition of sheer good luck.

³I follow Castenada's convention of adding an asterisk to pronouns corresponding to de se belief.

There is a kind of probabilistic incoherency in believing both that p and that there is a low probability that one's belief that p is true. Rationality demands a match between one's lower-order and higher-order probabilistic judgments; a match that Brian Skyrms labeled *Miller's principle*.^[12] Following Miller's principle involves transcending an unjustified egocentricity, treating oneself as a typical member of one's kind (in the absence of evidence to the contrary). Since internal justification is such an important part of our cognitive design plan and the respect for defeaters an important aspect of internal justification, knowledge is impossible in the presence of unneutralized defeaters.

Proposition 1 *S knows that p only if every prima facie rebutting or undercutting defeaters of p for S is neutralized for S.*

4.2 Why Immunity to Defeat is Needed

However, knowledge requires more than merely the absence of actual defeaters. One must be in a state of belief that is in certain respects *immune* to defeat. If S believes that p merely because he has not inferred certain propositions that one would be fully justified in inferring, then S 's continuing to believe p is simply a matter of dumb luck, which is incompatible with knowledge. Moreover, one must be, at least to some extent, resistant to misinformation and misleading suggestions that would undermine one's internal justification. If the mere suggestion that S might not be warranted would be sufficient to defeat his belief that p , then S 's belief that p doesn't have the kind of stability and robustness that is required for knowledge.

Proposition 2 *S* knows that *p* only if *S* is in a position to believe that *p* with internal justification.

Proposition 3 *S* is in a position to believe with internal justification that *p* only if *S* is in a position to believe with internal justification that it is very unlikely that *q* is true, for every available *q* that would, if believed, be an unneutralized defeater for *S* of *p*.

In particular, in order to know that *p*, one must be in a position to disbelieve with justification any possible *undercutting* defeater of *p*. As Plantinga has argued, one's belief that *p* is undercut if the probability that one's belief that *p* is warranted is either low or *inscrutable*. If a reasonable doubt about the epistemic status of one's belief that *p* can be raised, and one is not in a position to judge with reason that the probability of a failure of warrant is low, one would be unjustified to persist in believing that *p*. Suppose, to use Plantinga's example, that one has conflicting evidence that is hard to assess concerning whether or not the lighting on the widget factory floor is tinted. Even in the absence of positive belief that the lighting is tinted, one's putative knowledge about the color of the pink-looking widgets is undercut so long as one cannot justifiably assign a low probability to the condition of abnormal lighting.

An *available* defeater is one that *S* could think of and understand. A proposition that is simply beyond *S*'s ken cannot be, even potentially, a defeater of any of *S*'s belief. This has the somewhat paradoxical consequence that more sophisticated thinkers are, in certain, more vulnerable to epistemic defeat than are less sophisticated ones. Although paradoxical, this is nonetheless true, since

with greater sophistication comes greater sensitivity to certain kinds of reasons, which can in turn in some circumstances weaken rather than strengthen one's grip on the truth.⁴

By *being in a position* to believe p with justification, I mean that one can arrive (through the competent exercise of one's inferential capacities) in a stable noetic state in which one believes that p and in which one has access to exactly the same evidence as one has in one's actual noetic state. The *noetic state* of a person includes everything about him that is relevant to the internal justification of his beliefs: what the person believes, together with all of the relevant perceptual and mnemonic appearances. A noetic state is *stable* if it contains no beliefs that would be defeated by further competent ratiocination.

Proposition 4 *S is in a position to believe with internal justification that p only if S is in a position to believe with internal justification that it is highly likely that his* belief that p is warranted, since S is a member of contemporary society, for whom undercutting defeaters (of the form: his* belief that p is unwarranted because not caused in the right way) are available.*

This proposition doesn't give us closure of knowledge under the presupposition-relation. If one knows that p and the fact that q is a necessary condition of one's knowing that p , it does not follow that one knows that q . I can know that I am seeing a chair without knowing that I'm not a brain in a vat, even if my knowing that I'm seeing a chair presupposes that I'm not a brain in a vat.

⁴Thanks to Michael Rea for pointing out to me the necessity of some sort of availability condition on potential defeaters.

The immunity-to-defeat condition falls short of the Timothy Williamson's *KK* principle [13] in three ways: (i) actual belief is not required: one must merely be in a position to believe something, (ii) it is not necessary that one be in a position to *know* anything but merely to be justified in believing it, and (iii) what one must be in a position to believe is not that one is warranted, but only that it is *highly likely* that one is warranted.

Thus, it's enough if I can rationally assign a very low probability to the failure of a presupposition of knowledge. To know that I am seeing a chair, I must be in a position to judge it very unlikely that I'm a brain in a vat. I needn't *know*, or even believe, that I'm not. The probabilities involved are subjective in nature but correspond to the probability judgments of an ideal, rational agent. In order to be in a position to believe that there is a high probability that I am not a brain in a vat, I must believe that the world has been so constituted as to make the objective probability of my having become a brain in a vat quite low. Otherwise, the possibility that I might be a brain in a vat would have an inscrutable probability, and so would be a successful undercutting defeater.

In the case of empirical knowledge, one must be justified in believing with a high degree of rational confidence that any of the situations making up one's knowledge-net – one's belief-states that constitute ordinary empirical knowledge and any of the epistemically mediating events (events belonging to the causal chain that connects those belief-states with their objects, or to the common cause of the states and their objects) – are caused in an epistemically appropriate, normal way. Let's call the set of *S*'s beliefs about his* knowledge-net

S's *conception* of his* knowledge-net. Here we encounter another paradox: the richer and more detailed is *S*'s conception of his* knowledge-net, the greater is his potential vulnerability to undercutting defeaters. For each kind *K* such that *S*'s conception of his* knowledge-net includes the belief that some of the situations in the net are of kind *K*, *S* must be in a position to judge that it is very likely that situations of type *K* are caused (and caused in an epistemically appropriate way, given their putative role in *S*'s net).

4.3 Strong vs. Weak A Priori Justification

There are at least two categories of beliefs that must be distinguished from our ordinary empirical beliefs: beliefs that are justified in a *strongly a priori* way, and self-verifying beliefs.

There are two kinds of a priori justification: weak and strong. Weak a priori justification requires the occurrence of a certain kind of non-sensual experience: a cognitive or intuitive experience, in which a proposition somehow presents itself to the mind as apparently true, or a felt inclination, by which holding the belief in question seems natural or normal. In contrast, a belief is justified in a strongly a priori way if the belief is justified without reference to any kind of experience or inclination whatsoever, whether sensual or purely intellectual.

Weakly a priori knowledge would seem to be subject to undercutting defeat along lines very similar to those that pertain to ordinary empirical beliefs. If I were to believe that one of my beliefs that I took to be justified in a weakly a priori way was actually uncaused, or that the cognitive experience or inclination

on which that belief was based was uncaused, this would seem to constitute an undercutting defeater. This would apply also to those beliefs that are justified because we have a natural inclination to have them, of the kind that Thomas Reid described. If, in my estimation, this inclination to believe flows not from a human nature designed by God or by natural selection to grasp the truth reliably but is simply uncaused, then I would have a defeater for the Reidian justification of such a belief. In the case of intellectual experience, it may be unclear how exactly our intellectual intuitions are connected to the situations which are their subject matter, but it seems clear that an uncaused intuition has not been formed in a normal or reliable way. Hence, it is only the *strongly* a priori beliefs that are intrinsically immune to such defeat.

4.4 Self-Verifying Beliefs

There is another category of belief intrinsically immune to defeat through doubts about the scope of causation: self-verifying beliefs. A belief is self-verifying if its occurrence satisfies its own truth-conditions (like believing that I am believing something). I will also exclude beliefs that are causally upstream of their truth-conditions, such as beliefs about what I will do or ought to do. It seems plausible to suppose that I can know what I will do without making any assumptions about the causes of that belief, so long as I believe that it will have its usual effects.

Beliefs about one's current phenomenal state or about the contents of one's current thoughts fall into a similar category. It's reasonable to suppose that

these beliefs in some way incorporate their objects, in such a way that the belief-state as a whole includes the truthmaker for the proposition believed. Such reflective, introspective beliefs are thus also immune to defeat due to doubts about their causes.

4.5 The Principal Argument

To summarize what I have claimed so far, ordinary empirical knowledge is subject to the condition that it be immune to undercutting defeaters. In particular, we must be in a position to believe that it is highly likely that our ordinary empirical knowledge is caused, and this latter belief must not depend for its justification on the empirical knowledge whose causal status is in question. If such circular immunity were allowed, few empirical beliefs would ever be vulnerable to undercutting defeaters, since we could always appeal to the belief itself, along with similar beliefs, to justify a belief in the reliability of the process leading to it. For example, to return to the widget factory, if circularity were permitted, I could justify my belief that my perceptual beliefs were reliably (and therefore probably normally) caused by simply relying on my perceptually formed beliefs in the pinkness of the observed widgets to verify that my perceptually formed beliefs about the widgets are all true. To avoid such vicious circularity, we must posit a kind of asymmetric dependence between the actual justification of an empirical belief and the potential justification of a hypothetical belief in its probable warrant.

This asymmetric dependency relation should not be confused with a differ-

ent, distinct relation: that of believing one proposition *on the basis of* one's belief in another proposition, *inferring* one belief from another, or, conversely, one belief's providing one with *evidence* for a second belief. However, we should not assume that this evidential/inferential relation is the only form of justification-dependency. Providing immunity to defeat is a distinct form of dependency.

Proposition 5 *If S's belief that p is not strongly a priori justified or self-verifying (i.e., if S's belief that p is an ordinary empirical belief), then S is in a position to believe with internal justification that p only if S is potentially in a position to believe with internal justification that it is highly probable that his* belief in p is warranted (formed by a normal and alethically reliable process), and in such a way that S's belief that p would depend for its internal justification on the justification of the latter belief (i.e., of the belief that it is highly probable that his* belief in p is warranted).*

Proposition 6 *It is evident that (with the possible exception of strongly a priori justified beliefs and self-verifying beliefs) any belief that is uncaused or whose epistemic grounds are uncaused is not warranted (because such a belief is not then formed by a normal or alethically reliable process). Moreover, the proposition that some or all of his* beliefs are uncaused is available to S.*

Lemma 1 *If S's belief that p is an ordinary empirical belief, then: S knows that p only if S is in a position to believe with internal justification that it is highly probable that his* belief in p and its grounds are caused, and in such a way that S's belief that p would depend for its justification on the justification*

for the latter belief.

Because both the evidence available to us and our inferential capacities are finite, there can be no infinite regresses involving the justification relation between potential beliefs. Moreover, the justification relation is transitive and irreflexive, ruling out any justificatory circles.

Proposition 7 *Let R be a relation whose range is the set of propositions belief in which S is in a position to be internally justified in any having, and let R hold between two propositions p and q just in case S is potentially in a position to be in a noetic state in which S 's belief that q depended for its internal justification on S 's belief that p . Then R is a partial well-ordering (well-founded,⁵ transitive and irreflexive).*

Arguably, proposition 6 is too strong, since it may be that the justificatory relationship between two potential beliefs (say, a belief that p and a belief that q) might be indeterminate, in that there could be two stable noetic states available to S , n and n' , of such a kind that p depends on q in n , and q depends on p in n' . However, I am confident that the following lemma (Lemma 2) would still follow from a suitably weakened version of proposition 6. Proposition 6 could be reformulated in terms of the dependency relation that holds within each noetic state available to S . Even if S could muster some empirical evidence for the thesis that his* belief that p and its grounds are caused, those empirical beliefs would themselves require further beliefs about their causal history in order to

⁵A relation is]em well-founded when every non-empty subset has at least one minimal element. This rules out both cycles and infinite regresses.

secure their immunity from defeat. Eventually, S 's empirical resources must be exhausted, culminating in beliefs whose immunity from defeat is intrinsic (i.e., either self-verifying or strongly a priori justified beliefs).

Thus, since an ordinary empirical belief always depends on the potential justification of belief in the causation of that belief, and since the justification relation on such beliefs and potential beliefs is well-founded and asymmetric, it follows that empirical knowledge must rest ultimately on a belief in the scope of causation that is intrinsically immune to defeat involving doubts about causation, that is, on a belief that is not itself an ordinary empirical belief.

Lemma 2 *If S 's belief that p is an ordinary empirical belief, then: S knows that p only if there is a noetic state n and a proposition q of such a kind that (i) S is in a position to be in n , (ii) in state n , S 's belief that p depended for its internal justification on S 's belief that it is highly likely that his* belief that q is caused, and (iii) this latter belief would not depend for its internal justification on any ordinary empirical belief of S 's.*

However, the range of possible self-verifying beliefs is too limited to provide the relevant sort of beliefs about causation. Hence, ordinary empirical knowledge depends on the potentiality of strongly a priori justified beliefs in the causation of one's knowledge-net.

Proposition 8 *There is no noetic state n and proposition p such that S is in a position to be in state n and, in state n , S would believe that it is highly likely that his* ordinary empirical belief that p and its grounds are caused, and this*

latter belief would depend for its internal justification on any non-belief (i.e., perceptual or mnemonic state) or any self-verifying belief.

By process of elimination, since S 's belief that his* ordinary belief that p and its grounds are likely to be caused cannot be supported by any of S 's ordinary empirical beliefs, any of his self-verifying beliefs, or any of his epistemically relevant non-beliefs, S must be strongly a priori justified in believing that it very likely that this belief and its grounds are caused.

Lemma 3 *If S 's belief that p is an ordinary empirical belief, then: S knows that p only if there is a proposition q such that S is potentially in a position to be SAP (strongly a priori) justified in believing that his* ordinary empirical belief that q is caused, and in believing that the epistemic grounds of his* belief that q are also very likely to be caused.*

It is implausible to suppose that one could be SAP justified in believing something about the likely causal history of some ordinary empirical beliefs but not others. SAP justified immunity to undercutting defeat will either extend to the whole of S 's knowledge-net or to nothing at all. A SAP justified belief is one whose content is a necessarily true proposition⁶ belief in which is constitutive of rationality, in the sense that it is a necessary condition of any ordinary empirical knowledge of the kind possessed by human beings. Such SAP justified beliefs must be suitably general in scope, embracing all possible items of human knowledge.

⁶Or, at least, a proposition whose truth is necessitated by the existence of human knowledge.

Proposition 9 *Necessarily, if S 's belief that p is an ordinary empirical belief, then S is potentially in a position to be SAP justified in believing that it is highly likely that his* belief that p and its grounds are caused only if S is in a position to be SAP justified in believing that it is highly likely that any of the situations in his* empirical knowledge-net are caused.*

Even if I am wrong about this, the the scope of our a priori knowledge of causation will be essentially unaffected. Suppose that SAP justification is *particularistic*: one can be justified in believing that it is likely that this or that belief (and its grounds) are caused, but never justified in believing the corresponding generalization about all of one's knowledge-net. Since there are no a priori limits to what can be thought of, one's knowledge-net could embrace the whole of the Cosmos, in which case one would be SAP justified in believing that the Cosmos as a whole has a cause.

Lemma 4 *Necessarily, if S 's belief that p is an ordinary empirical belief, then S knows that p only if S is in a position to be SAP justified in believing that it is highly likely that any of the situations in his* empirical knowledge-net (i.e., his ordinary empirical beliefs and their epistemic grounds) are caused.*

In order to be justified in believing that it is very likely that any situation of one's knowledge-net is caused, one must be able to apply some principle of general causation to a category that includes nearly all of that knowledge-net.

Definition 4 γ is a principle of general causation iff γ takes the form

It is nomologically impossible for a situation of type T to be actual in the absence of a cause.

For such a principle γ , T is γ 's *range of application*.

In order for S to be SAP justified in believing that it is highly likely that each of the situations in his* knowledge-net is caused, there must be some type T such that S is SAP justified that every situation in his* knowledge-net belongs to T , and that the probability that a given event has a cause, conditional on its being an actual member of T , is very high. This conditional probability is high only if S can judge the probability of the disjunction of the following three conditions to be very high:

1. If a possible event e is of type T , then there is a non-negligible prior probability there actually exists a potential cause of e 's occurrence.
2. If an event e is of type T , then the law-based probability of e 's occurring uncaused must be much lower than the probabilistic expectation of its occurring as the result of a cause.
3. It is nomologically impossible for an event of type T to occur in the absence of a cause.

However, it is clear that there is no type T that both self-evidently⁷ applies to every situation in S 's knowledge-net and that meets either condition 1 or

⁷By *self-evident*, I mean that S is justified in believing the proposition on the basis of SAP-justified and self-verifying beliefs alone.

condition 2. *S*'s knowledge-net includes situations of arbitrarily low prior probability, and, therefore, the prior probability of the occurrence of a cause of a situation in that net is also arbitrarily low. Hence, the rational probability of *T*'s meeting condition 1 must be very low. *S* must be justified in judging, for *each* token-situation in his knowledge-net, that the probability of *its* occurring without a cause is very low, no matter how low the prior probability is of the occurrence of that token or of its potential causes.

For the same reason, no matter how low the law-based probability of the uncaused occurrence of an event *e* might be (assuming that this probability is greater than zero), the probability of the occurrence of a cause of *e* could be lower still. Thus, condition 2 also has a vanishingly small probability of truth. It is, a priori, very unlikely that the laws of nature should be jury-rigged in such a way to make condition 2 come out as true, since this would require the probability of a token situation's uncaused occurrence being highly sensitive to the probability of the occurrence of one of its potential causes. Thus, the only way that *S* could be SAP justified in judging that the probability of the uncaused actuality of *each and every one* of the token situations in his* knowledge-net is low would be by judging the probability of condition 3 to be high.

If I'm wrong about conditions 2, then the right conclusion to reach would be that *S* is SAP justified in believing that his* knowledge-net falls within the range of some *qualified* principle of general causation:

Definition 5 *γ* is a qualified principle of general causation iff *γ* takes the following form:

The objective probabilities are of such a kind that, for every possible situation s of type T , the probability of s 's occurring uncaused is vanishingly low (so low that, no matter how unlikely the caused occurrence of s might be according to a possible noetic state, its uncaused occurrence is much more unlikely).

Proposition 10 *S is potentially in a position to be SAP justified in believing that it is highly likely that any of the situations in his* empirical knowledge-net are caused only if S is SAP justified in believing that it is very likely that there is some type T such that (i) some principle (or qualified principle) of general causation γ holds with T as its range of application, and (ii) it is self-evident to S ⁸ that nearly all of the situations in his* empirical knowledge-net fall within T .*

Since there are a small number of types that could meet both conditions (i) and (ii) of proposition 10, and since if types T_1, T_2, \dots, T_n each individually meet the pair of conditions, so does their conjunction (since both nomological necessity and self-evident justification are closed under conjunction), we can shift the relative positions of the probability operator and the quantification over types in proposition 10, resulting in proposition 11:

Proposition 11 *S is potentially in a position to be SAP justified in believing that it is highly likely that any of the situations in his* empirical knowledge-net are caused only if there is some type T such that S is SAP justified in believing that (i) it is very likely that it is nomologically impossible for situations*

⁸That is, S 's belief is justified on the basis of a combination of self-verifying and SAP justified beliefs.

of type T to be actual in the absence of a cause, and (ii) it is self-evident to S that nearly all of the situations in his* empirical knowledge-net fall within T .

From lemma 4 and proposition 11, the main result follows.

Theorem 1 *If S 's belief that p is an ordinary empirical belief, then S knows that p only if there is some type T such that S is SAP justified in believing that (i) it is very likely that it is nomologically impossible for situations of type T to be actual in the absence of a cause, and (ii) it is self-evident to S that nearly all of the situations in his* empirical knowledge-net fall within T .*

4.6 Consistency of this Requirement with a Reidian, Common Sense Epistemology

Unlike Descartes, Hume, Kant and other early modern epistemologists, I am not assuming that strongly a priori justified beliefs are indubitable, incorrigible, or associated with any special phenomenal quality (like Descartes' clear and distinct ideas). I have argued for the claim that we are SAP justified in believing a causal principle on Moorean grounds: assuming that we do in fact have empirical knowledge, and arguing on that basis that we must be SAP justified in believing the causal principle. I have not argued that we must actually believe the causal principle in order to have empirical knowledge, merely that we must be in a position to do so with with SAP justification.

My conception of SAP justification is something like this: to be SAP justified in believing that p , the disposition to believe p must somehow be constitutive

of rationality itself. That fact that it is so may itself be a kind of a posteriori discovery. Thus, I would reject the assumption that if one is SAP justified in believing that p , one must be SAP justified in believing that one is so justified.

It is surprising, to say the least, to find that there must be some SAP justified beliefs. I myself was surprised by this result: I had defended (in *Realism Regained*[5]) the thesis that all of our knowledge, even our knowledge of the laws of logic, depended on the right kind of causal connection to the relevant facts. I now see that I was mistaken: if we have any knowledge at all, some of that knowledge must be absolutely immune to undercutting defeaters. An SAP justified belief is such that, even if I came to believe that the belief in question had been planted in my mind by a malicious Cartesian demon, I would still be rationally justified in persisting in my belief. I should, in such a case, merely conclude that the demon had, on this occasion, inadvertently supplied me with genuine knowledge.

5 Epistemologically Acceptable Exceptions to the Universality of Causation

Firstly, if one is SAP justified in believing that one's knowledge-net falls within the range of application of some PGC (principle of general causation), then it must be metaphysically necessary that any knowledge-net of any person with humanoid consciousness (one who shares with us the same sort of cognitive make-up, the same sort of self-verifying beliefs about his* own mental state,

and a similar conception of his* knowledge-net) falls within that same range.

Secondly, the PGC must be crafted in such a way that it is plausible to suppose that it is *self-evident* that its range of application encompasses our knowledge-nets. Our knowledge of the past and of the extent of the universe is entirely a posteriori, as can be seen through the thought-experiment of a humanoid creature with no memories of the past, no traces or testimony of past events, and no sensory presentations or abductive evidence of spatially remote events.

Thirdly, an epistemologically acceptable PGC must be one whose range of application specifies a set of *intrinsic properties* of situations. It is highly implausible that the nomological possibility of a situation's being caused or uncaused has anything to do with that situation's extrinsic properties, including what sorts of effects it might have. Moreover, it is only the intrinsic character of our belief states (their content and associated phenomenological qualities) that are self-evident to us, and it is only the intrinsic natures of things (conferring causal powers) that can be the object of abductive inference.

Fourthly, the boundaries of the range of application of the PGC must be non-arbitrary and metrically isolated, if belief in the PGC is to be SAP justified. If there are cases that are close to, but not within, the range of application of the PGC, then Timothy Williamson's margin-of-error principle would prevent our being justified in applying the principle at its margins. However, this contradicts Theorem 1, which states that we are SAP justified in our confidence in the PGC.

Fifthly, the range of application of the PGC must be closed under proper

parthood. That is, if x is a member of the range of application T , and y is a proper part of x , then y must also be a part of T . Conversely, if we are not in a position to presume that each of a thing's parts have a cause, then we are not in a position to presume that the whole has a cause, either. It might be thought that there are cases in which a complex situation should be presumed to have a cause, independently of any such presumption about its parts, because the complex situation incorporates many coincidences that demand an explanation in terms of a common cause. However, this overlooks the fact that such an inference to the best (causal) explanation is plausible only when the phenomenon to be explained is already supposed to have a cause. Occam's razor gives us grounds for preferring simpler causes to more complex ones, but it gives us no reason to prefer any cause (even a very simple one) to no cause at all.

Finally, an epistemically acceptable PGC must be sensitive to the fact that human cognition includes an open-ended, highly general form of *abductive reasoning*. In order for a situation to constitute possible abductive evidence for a further conclusion about its causes, only two things are required: (1) the situation must have an intrinsic character that is compatible with its having causes, and (2) the situation must have an intrinsic character of a kind that kind figure in possible causal laws, which entails that the character is multiply realizable.

In light of these criteria, the following would seem to be a model of an epistemologically acceptable principle of general causation:

Principle 1 *It is (at least) nomologically necessary that if x is composed of some y 's, and for each y among the y 's, it is metaphysically possible that there*

exist a situation that approximately duplicates y and has a cause, then it is highly likely that x itself has a cause.

The range of Principle 1 is clearly determined by a situation's intrinsic characteristics, and it is closed under parthood. In addition, Principle 1's range of application has non-arbitrary boundaries: anything that is very similar to a member of its range also belongs to the range. Principle 1 also respects the open-endedness of abduction, since any phenomenon that could possibly justify an inference to a best causal explanation would have to have an intrinsic character compatible with its having a cause. Moreover, the range of application of Principle 1 is wide enough to encompass the entire knowledge-net of any humanoid creature, and to do so self-evidently. Principle 1 is at least a paradigm of an epistemically acceptable version of the PGC.

5.1 Unacceptable Candidates

As I argued above (in section 3), any acceptable principle of general causation (PGC) must be restricted in its scopes. Given Theorem 1, we are now in a position to assess various candidate restrictions for their compatibility with securing immunity to defeat for our empirical knowledge. The crucial question becomes: can we be SAP justified in believing that nearly all of our knowledge-net falls within the range of application of the proposed PGC?

Let's consider some of the candidate restrictions from section 3.

1. All non-first situations have causes.

How could I be SAP justified in believing that my current belief-state is a non-first situation? In order to do so, I would have to know that there were situations that preceded my current belief-state in time, but my knowledge of the past consists entirely in ordinary empirical beliefs (including memory and testimony), all of which presuppose (as I have argued) belief in the applicability of the causal principle to my current belief-state. Hence, we cannot have a non-circular justification of immunity to defeat, relying on this version of the PGC.

2. All events that don't occur at a first moment in time have causes.

In the same way, it is impossible for me to be SAP justified in believing that my current belief-state is not occurring at a first moment of time. All of my knowledge about the past, even that there were past moments, is dependent on that ordinary empirical knowledge whose immunity from defeat is at issue. It is certainly metaphysically possible for a person with consciousness like mine to exist in a first moment of time.

Moreover, even if it were possible to be SAP justified in believing that my *beliefs* don't occur at a first moment in time, there would be no way to generalize this belief to all of the situations in my knowledge-net, many of which occur in the remote past.

Finally, this principle fails to have a range of application that refers only to the *intrinsic character* of situations. If any sort of situation could occur at a first moment of time and thus lack a cause, then any sort of situation could occur uncaused. This principle reduces to the trivial assertion that situations

have causes, unless they don't.

3. All events with finite temporal duration have causes.

Just as I cannot be SAP justified in believing that my current belief-state is not occurring at a first moment in time, so I cannot be SAP justified in believing that that same state has not persisted for an infinite period of time. I know that my current belief-state has endured only for a finite period of time only because I have ordinary empirical knowledge of events that preceded it in time.

A fortiori, it is impossible for me to be SAP justified that my entire knowledge-net has only a finite temporal duration.

4. All events not including infinite causal regresses have causes.

How do I know that the processes internal to my knowledge-net don't include infinite causal regresses? If time is dense, then infinite causal regresses may be quite common throughout the physical world, as Alexander Pruss argued in "The Hume-Edwards Principle".[11] Consider Pruss's example of the flight of a cannonball. If time is dense, then there are infinitely many dynamic states of the cannonball occurring after the firing of the cannon. For example, there are states one second after the firing, one-half second after, one-quarter second after, and so on. Each of these states is causally prior to its successors, and in each case the energy and momentum possessed by the cannonball in that state is sufficient to explain its subsequent trajectory.

Even if it is in fact false that such infinite regresses occur within our knowledge-net (if, for example, it turns out that time is discrete), our knowledge of such a

fact could hardly be strongly a priori.

5. All events that are both extremely simple (have extremely low entropy) and cosmic in scope have causes.

It is not at all obvious that my current belief state isn't simple, and I am not SAP justified in believing that my current state doesn't exhaust the cosmos. My belief that there are many currently existing states outside myself is dependent on my ordinary empirical knowledge.

Much depends on how the notion of *extreme simplicity* is cashed out. If we interpret simplicity in Neo-Platonic, Plotinian terms, as a kind of metaphysically absolute simplicity, then this version of the PGC will work epistemologically, but it will also be strong enough to support a theistic conclusion. Alternatively, if we interpret *simplicity* in terms of thermodynamic entropy, then it would be incredible to suppose that I have *any* SAP justified beliefs about the entropy of various states.

Furthermore, this version of the PGC picks out events by virtue of one of their extrinsic features: being cosmic in scale (i.e., the non-existence of contemporaneous disjoint events). Such features couldn't plausibly be relevant to the chances of an event's having a cause. Nor is it plausible that I am SAP justified in believing that my belief state is not cosmic in scale.

6. All situations that could (de re) be caused have causes.

There is no way that I could be SAP justified in believing that my current belief-state could (de re) have a cause. If we suppose that each thing necessitates

de re its particular origins, then it would follow that any uncaused state is essentially uncaused. This version of the PGC would then collapse into the trivial principle that whatever has a cause has a cause.

7. All situations that are either ordinary empirical beliefs of mine or that mediate causally between those beliefs and their objects (or between those beliefs and causes of their objects) have causes.

There would, of course, be no problem about my being SAP justified in believing that all the situations in my knowledge-net fall within the scope of application of this principle, since this is a trivial logical truth. However, this is an unacceptable PGC because it does not specify its range of application in terms of the *intrinsic character* of situations. It is surely not justified (not to mention SAP justified) to believe that whether or not a situation is (by virtue of its causal connections with my behavior) among my beliefs or has one of my beliefs among its effects is in any way relevant to the probability that the situation has a cause. Nor would it be justified to believe that it is a metaphysical necessity that these extrinsic features are correlated with some intrinsic character that nomologically entails having a cause.

One could perhaps argue that human belief-states have the sort of intrinsic character that demands, not only that they probably have causes, but that they are probably caused in a particular way (by a long process of natural selection, for example), a way that implies that it is likely that they are causally connected to their objects. It may be thought, for instance, that the functional unity and complexity of human thought entails the likelihood of such an explanation.

However, as I argued above, this sort of inference to the best (causal) explanation presupposes that some more general PGC applies to the situations in question, taken one at a time: it cannot be used as the ground for a more restricted PGC that applies only to the whole. Inference to the best explanation gives us reason to choose the simpler (ontologically weaker) of two causal explanations of a phenomenon: it does not give us a reason to prefer a causal explanation over the hypothesis that the phenomenon was uncaused. If the hypothesis that a situation was uncaused is a viable one, then Occam's razor would always give us reason to prefer that hypothesis over *any* causal explanation, since the utter absence of any cause is simpler (ontologically leaner) than any hypothesis of a cause could be. Thus, unless the particular situations making up our knowledge-net have some intrinsic character that we are SAP justified in believing to entail a high probability of causation, we are in no position to believe that the knowledge-net as a whole has such a cause.

6 Conclusion

I have argued that it is surprisingly difficult to avoid both the Scylla of skepticism and the Charybdis of apriorism about the scope of causation. This is, to be sure, the first rather than the last word on the subject. There is a great deal more work to be done on the nature of immunity to defeat, on the nature of the related dependency relation between propositions, and on the possible scope of strongly a priori beliefs. However, it seems clear that some sort of a

priori commitment to the generality of causation is required to secure warrant for our ordinary empirical beliefs, even though many of those empirical beliefs are basic (in Plantinga's sense) rather than inferred.

References

- [1] Ernest W. Adams. *The Logic of Conditionals: An Application of Probability to Deductive Logic*. D. Reidel, Dordrecht, 1975.
- [2] Ernest W. Adams. The logic of high probability. *Journal of Philosophical Logic*, 15:255–279, 1986.
- [3] Edmund Gettier. Is justified true belief knowledge? *Analysis*, 23:121–123, 1963.
- [4] Robert C. Koons. A new look at the cosmological argument. *American Philosophical Quarterly*, 34(193-211), 1997.
- [5] Robert C. Koons. *Realism Regained: An Exact Theory of Causation, Teleology and the Mind*. Oxford University Press, New York, 2000.
- [6] Robert C. Koons. Defeasible reasoning, special pleading and the cosmological argument: A reply to oppy. *Faith and Philosophy*, 18(192-203), 2001.
- [7] Graham Oppy. Koons' [sic] cosmological argument. *Faith and Philosophy*, 16:378–389, 1999.

- [8] Alvin Plantinga. *Warrant and Proper Function*. Oxford University Press, New York, 1993.
- [9] John Pollock. Defeasible reasoning. *Cognitive Science*, 11:481–518, 1987.
- [10] Huw Price. *Time's Arrow and Archimedes' Point: New Directions for the Physics of Time*. Oxford University Press, New York, 1996.
- [11] Alexander R. Pruss. The Hume-Edwards principle and the cosmological argument. *International Journal for the Philosophy of Religion*, 43:149–165, 1998.
- [12] Brian Skyrms. *Pragmatics and Empiricism*. Yale University Press, New Haven, Conn., 1984.
- [13] Timothy Williamson. *Knowledge and its Limits*. Oxford University Press, New York, 2000.