

Chapter Five

Can Skepticism Be Refuted?

A standard form of skepticism rests on the premise that the choice between our ordinary beliefs and skeptical hypotheses is evidentially underdetermined. To make a case against this premise, Jonathan Vogel invokes inference-to-the-best-explanation and considers what he calls the “isomorphic hypothesis,” which matches the cause and effect relationships of the real world hypothesis, according to which things are by and large what we take them to be. The isomorphic hypothesis’s explanatory power seems to be no less than that of the real world hypothesis. On what grounds, then, can we prefer the latter over the former? Vogel argues that the isomorphic skeptical hypothesis does not deepen our understanding of our perceptual experiences. The isomorphic hypothesis is more complex than the real world hypothesis because it replaces genuine shapes with pseudo-shapes and genuine locations with pseudo-locations. This increased complexity is not offset by a commensurate increase in explanatory scope and depth. Hence, Vogel suggests, we may prefer the real world hypothesis over the skeptical isomorphic hypothesis and thus reject the evidential underdetermination premise on which skepticism relies.

In his response essay, Richard Fumerton explores the challenge of refuting skepticism within the traditional framework of inferential internalism, according to which knowing propositions about the external world requires (a) that we have evidence E making such propositions probable and (b) that we know E makes such propositions probable. For inferential internalists to avoid skepticism, they must meet two serious challenges. First, to avoid an infinite regress, they must show that we can have direct knowledge of what makes what probable. Second, they must identify what body of evidence enables us to know that skeptical hypotheses are false. Does inference-to-the-best-explanation help with the latter challenge? Fumerton expresses doubt about that. Skeptical hypotheses might provide simpler explanations than their non-skeptical counterparts. Moreover, the skeptics have every right to question that inference-to-the-best-explanation is a guide to truth.

The Refutation of Skepticism

Jonathan Vogel

I

We take ourselves to know a lot about the world, and it would be profoundly disturbing if we didn't. To deny that we enjoy such knowledge is to endorse *skepticism* about the external world.¹ Skepticism is philosophically important because a gripping line of thought seems to show that it is correct.

The argument that supports skepticism is one of the most famous in the history of philosophy.² It turns on the possibility that we might be victims of some kind of massive sensory deception. Consider two situations. In one, everything is normal, and you see a bridge. In the other, you are the subject of an awful experiment. Your brain is isolated in a laboratory vat and fed completely delusory sensory inputs. These inputs make it appear to you as though there were a bridge before you, even though there isn't. It may well seem that you have no way of knowing which of these two situations you're in. So, in particular, you don't know that there is really a bridge before you. This line of thinking generalizes. If you don't know that you're not the victim of massive sensory deception, then you're unable to know virtually anything at all about the world. Skepticism prevails.

That is the gist of the "deceiver argument" (as I shall call it), but there are various details that need to be taken into account. I think we can understand better what is at stake by comparing the deceiver argument with an ordinary case in which someone fails to know. Imagine that you're in the kitchen, and the toaster suddenly stops working. It could be that the toaster itself has burned out. However, it could also be that the toaster is all right, but a fuse has blown instead. You are, then, faced with two competing hypotheses as to what caused the toaster to shut off.³ If you have no further relevant information, you have no basis for accepting one hypothesis over the other. Any choice on your part would be arbitrary. I will say that, in such a situation, your choice of hypothesis is *underdetermined*. If, despite such underdetermination, you guess that the toaster shut off because the fuse blew, you wouldn't *know* that – even if, luckily, your guess turned out to be correct. This sort of example supports the view that knowledge is governed by the following general principle:

Underdetermination principle (UP): If q is a competitor to p , then one can know p only if one can non-arbitrarily reject q .

Rejection of q would be arbitrary in the relevant sense just in case q is, from an epistemic standpoint, no less worthy of belief than p . So, we have:

Underdetermination principle (UP, alternate version): If q is a competitor to p , then a subject S can know p only if p has more epistemic merit (for S) than q .

What factors add to or subtract from epistemic merit is a crucial, but controversial, matter. I address it below.

With these points in hand, we can set out the deceiver argument more rigorously, as follows.

- (1) Consider any proposition *m* about the world I ordinarily believe (hereafter, “mundane propositions”). In order to know that *m*, my belief that *m* must not be underdetermined.
- (2) My belief that *m* is underdetermined.
- (3) Therefore, I don’t know that *m*.

Premise (1) may be seen as a straightforward application of the underdetermination principle. According to that principle, underdetermination is inimical to knowledge. Premise (1) makes the point that we fail to know any mundane proposition for which there is a competitor of equal epistemic merit. Premise (2) is the claim that all mundane propositions face a competitor of just that sort. The competitor is that one is the victim of massive sensory deception, as described above (I will call this the “skeptical hypothesis”). The conclusion (3) is that we lack knowledge of mundane propositions in general, which is to say that skepticism about the external world holds. Now, the deceiver argument is logically valid. To refute it, we would have to successfully challenge one of its premises. Premise (1), and the underdetermination principle that underwrites it, seem unassailable.⁴ So, a satisfactory response to the deceiver argument will have to show that premise (2) is false, or at least not adequately supported.

But satisfactory to whom? The deceiver argument explicitly or tacitly depends on various epistemic principles. One is the underdetermination principle. In determining whether underdetermination exists, other principles will be brought to bear. Now, I take it that we are committed to a body of epistemic principles that govern what we count as knowledge, justified belief, and the like. The deceiver argument may be construed as (putatively) proceeding from just those principles. Thus understood, the argument is an attempt to show that, by our own lights, we lack the knowledge of the world we think we have. I will call this position *domestic skepticism*. Domestic skepticism is concessive, but dangerous. It is concessive, in that it doesn’t contest the legitimacy of the epistemic principles we embrace and employ. But it is also dangerous, in the sense that it would be deeply unsettling, or worse, if we have no knowledge of the world, according to our own accepted view of what knowledge is and what it requires. On the other hand, if it emerges that our epistemic principles don’t have that consequence after all, domestic skepticism has been refuted.⁵

My primary concern here will be with domestic skepticism. However, there are other forms skeptical thinking might take. Someone might contest not only our ordinary judgment that we have knowledge of the world, but also the legitimacy of the principles on which we rely in making that judgment. I will call this more extreme position “exotic” skepticism. It won’t matter to an exotic skeptic that our knowledge claims accord with our epistemic principles, since the latter themselves are supposed to be questionable in some way. Thus, a reply that answers the domestic skeptic will seem to the exotic skeptic like our pointlessly patting ourselves on the back. I will return to the difference between domestic and exotic skepticism in section III.

II

The case for skepticism depends on the status of premise (2) of the deceiver argument, and the status of that premise depends on whether mundane propositions have more epistemic merit than the skeptical competitors they face. Philosophical opinion varies widely and sharply on this point. One can distinguish the following positions, among others. (a) We have no basis whatsoever for rejecting skeptical hypotheses, if such hypotheses really do compete for our acceptance. Skepticism is at least “conditionally correct,” and it can’t be refuted.⁶ (b) While no particular evidence counts against skeptical hypotheses, epistemic rationality permits us or requires us to reject such hypotheses out of hand. “A reasonable man does not have certain doubts,” as Wittgenstein said (Wittgenstein, 1972, Remark no. 220). (c) What licenses us in maintaining belief in mundane propositions, and rejecting skeptical hypotheses, isn’t a special anti-skeptical epistemic principle, but a broader principle of methodological conservatism. Methodological conservatism is, roughly, the doctrine that we are entitled to maintain beliefs we already have, all other things being equal. Since we have already accepted mundane propositions, rather than any skeptical competitor, the former have a kind of epistemic merit the latter don’t. (d) Experience itself provides immediate justification for the acceptance of mundane propositions. On one way of framing this approach, it is a fundamental epistemic principle that if it appears to S that F, S is justified in believing that F. Insofar as experience justifies believing F, F is epistemically more meritorious than the hypothesis of massive sensory deception, and there is no underdetermination (see Pryor, 2000). (e) It has been argued that skeptical hypotheses are “dialectically” or “pragmatically” self-refuting, or even logically inconsistent. If this is so, and our mundane beliefs aren’t similarly defective, then our mundane beliefs are certainly superior to their competitors.⁷

I regard these views as worthy of serious consideration, but, in my judgment, all of them are ultimately untenable.⁸ The alternative I favor may be outlined as follows: when one is choosing between competing candidates for belief A and B, one has good reason to accept A rather than B if A provides a better explanation of a relevant body of facts than B does. That is, I uphold the legitimacy of what is known as *inference to the best explanation* (hereafter, “IBE”). Skeptical hypotheses are less successful than our mundane beliefs in explaining a relevant body of facts, namely, facts about our mental lives.⁹ Because of this disparity, we have reason to favor mundane propositions over their skeptical competitors, and the choice between them isn’t underdetermined after all. Therefore, premise (2) of the deceiver argument is false, and skepticism stands refuted.¹⁰ This approach faces some significant obstacles, however. There is unclarity about what counts towards explanatory goodness, and about how such goodness could be assessed or compared – there is even substantial disagreement as to what constitutes an explanation at all.¹¹ I certainly don’t propose to resolve these difficulties here. Rather, in what follows, I will make some specific assessments of explanatory goodness that I take to be plausible in their own right.

A further complication is that the skeptical hypothesis that one is the victim of massive sensory deception could be developed in any number of ways. Presumably, these will vary with respect to their degree of explanatory goodness. Consider what I will call the “minimal skeptical hypothesis” (MSH). The content of the MSH is just

that your experience is caused in a delusory manner, and no more: with respect to every Z, if it appears to you that Z, then something causes it to appear to you *falsely* that Z. (If you like, you can embellish the MSH a little by adding that the cause is an evil demon or the computer of a mad neuroscientist.) It seems to me that the MSH has little, if any, explanatory power. The putative explanations it offers are impoverished and ad hoc. According to the MSH, it appears to you that Z because something causes it to appear to you falsely that Z. That is like explaining that you fell asleep because of the action of something with a dormitive virtue – that is, like saying that you fell asleep because something caused you to fall asleep – which is practically no explanation at all.¹²

However, I have offered no reason to suppose that all skeptical hypotheses have to fail as badly as the MSH. The possibility apparently lies open that a skeptical hypothesis could posit a nexus of causes and regularities that is as cohesive and economical as the body of our ordinary beliefs about the world. Call this body of beliefs the “real world hypothesis” (RWH). In fact, it seems that the skeptic could take over the causal-explanatory *structure* of the RWH, but substitute within it reference to objects and properties other than the ones we take to be real. For example, suppose you move your hand so as to pet a cat, which purrs in response. The motion of your hand (H) causes your (V) visual experience as of a hand moving, and it causes (C) the cat to purr. The cat itself causes you to have (P) an auditory experience of a gentle rumbling. But now imagine that the brain-in-a-vat skeptical hypothesis were true, and that the computer has a file (H*) that stands in for your hand and another file (C*) that stands in for the cat. (H*), rather than a real hand, causes (V) and also activates (C*) the cat-file in a particular way. The cat-file, in turn, causes (P) your auditory experience of a purring sound.¹³ I will call a skeptical competitor of this sort the “isomorphic skeptical hypothesis” (ISH). The relationships among causes and effects according to the ISH match those of the RWH. To that extent, it seems, the explanations provided by the one are no better or worse than the explanations provided by the other. There is, then, no difference in explanatory success that favors the RWH over the ISH.

So far, we have considered two ways the skeptic might proceed. Either he deploys the MSH (or some variant of it) as a competitor to the RWH, or he puts forward a skeptical hypothesis that is structurally identical to the RWH. The MSH differs in structure from the RWH, but is explanatorily inferior to it. The ISH is supposed to have the same structure as the RWH, and, consequently, explain just as well as the RWH does. But it may seem that the inadequacies of the MSH needn't carry over to every other skeptical competitor that differs in structure from the RWH. So there could be some such competitor that equals the RWH in explanatory merit.

There are several points to make in response to this suggestion. First, a certain class of skeptical hypotheses (the MSH and its ilk) may be rejected on explanatory grounds. Formulations of the deceiver argument typically invoke skeptical hypotheses of just this kind, and those versions of the argument *are* refuted by the appeal to explanatory considerations. Second, the claim that, for some particular hypothesis H, there is a competitor that differs in structure from H, but equals H in explanatory merit, is by no means trivial and cannot be taken for granted. Adducing an equally good, structurally different competitor to a given hypothesis requires one to undertake substantive theory-construction, which is by no means guaranteed to succeed.¹⁴ Finally, as I will suggest below, the RWH enjoys certain explanatory advantages over

skeptical hypotheses. There is no reason to suppose that these advantages are offset by unspecified *disadvantages* attaching to the RWH, which a structurally different skeptical competitor will somehow manage to avoid. Hence, it is doubtful that one could cobble together a structurally different skeptical competitor to the RWH that will succeed in matching the RWH in explanatory merit.

To provide a setting for further argument, it will be helpful to say something more about the nature of our problem and its solution. As I see things, the underlying motivation for skepticism is that perception, the ultimate source of whatever empirical knowledge we have, is a causal relation. Our perceptual experiences are the effects of the world upon us. In general, causal relations are contingent, so there seems to be no bar in principle to anything causing a given effect. Various skeptical hypotheses serve to make us vividly aware of this point in the perceptual case. Our perceptual experiences could be caused by familiar sorts of things in the world, but those experiences could also be caused by radically different sorts of things (a computer, a demon). In the same way, it seems that specifying any wider structure of causal relations and regularities doesn't fix which properties and objects participate in that structure. Unfamiliar properties (e.g., a computer disk's magnetic patterns, a demon's mental states) could stand in for the familiar ones invoked by the RWH, without any additional loose ends or explanatory regularities. The result would be a skeptical hypothesis that has the same structural explanatory virtues as the RWH.

At the same time, though, there is a natural impression that a good explanation for why X behaves like something that is F is that X is, indeed, F. An explanation of why something that *isn't* F nevertheless behaves as though it were F seems bound to involve greater complications. However, someone with Humean scruples may well find such talk, as it stands, mysterious and misguided. She will think that X's being F by itself implies nothing about how X "behaves." X's being F doesn't determine what regularities X enters into, including whether X appears to us to be F or not.

The issues here are deep and hard, but I want to suggest that explanatory considerations can still do the requisite anti-skeptical work.¹⁵ Even though the ISH is supposed to share the structure of the RWH, it is meant to be a thoroughgoing competitor to it. That is, the ISH is required to explain the character of our experience as successfully as the RWH does, while positing a very different disposition of matter in space. Rather than there being ordinary arrangements of ordinary-shaped objects, there is just your brain, its vat, and the computer with which it interacts.¹⁶ Now, the ascription of shapes and locations to things does a great deal of explanatory work for us. The ISH has to achieve the same explanatory success by ascribing properties other than shape and location to some other kind of object. Call these alternative properties "pseudo-shapes" and "pseudo-locations." My proposal is that skeptical hypotheses are bound to suffer an explanatory disadvantage insofar as they have to make do with pseudo-shapes and pseudo-locations instead of genuine shapes and locations.

The disadvantage is a lack of simplicity. It is, in principle, a defect to proliferate explanatory apparatus without any commensurate increase in the depth or scope of what is explained. So, other things being equal, a hypothesis that invokes fewer explanatory regularities is preferable to a competitor that invokes more. I take it that we are able to recognize differences of this sort and choose between hypotheses accordingly. For example, we find a version of mechanics that provides a uniform account of terrestrial and celestial motion preferable to a mechanics that invokes one

set of laws governing the former and another set of laws governing the latter. Similarly, we accept that most people who seem comfortable really are so, rather than “super-spartans,” that is, people who are in great pain, yet have the desire and fortitude not to manifest it in any way.¹⁷ I believe that the RWH is more economical than the ISH in much the same way, and so the RWH may be favored over the ISH on explanatory grounds.

In previous work, I tried to establish that result as follows (Vogel, 1990). It is, I claimed, a necessary truth that two distinct objects can't be in the same place at the same time.¹⁸ This truth about location is deeply ingrained in our understanding of the world. If I walk two blocks north and then two blocks west, while you walk two blocks west and then two blocks north, we encounter the very same thing, whatever it may be, when we arrive. In the RWH, such facts are explained by the necessary truth that if X has the property of being located at a genuine location L, and Y is distinct from X, then Y doesn't have the property of being located at L. Now, where the RWH says that two objects don't have the same location at the same time, the ISH must say that two pseudo-objects don't have the same pseudo-location at the same time. But in so doing, the ISH must invoke an extra empirical regularity, since it is metaphysically possible for distinct objects to share any property other than location.¹⁹ To make this point more vivid, imagine that the nefarious computer in the skeptical scenario keeps track of the “location” of “your hand” by writing coordinates in the hand-file, and likewise for the pseudo-location of other pseudo-objects. The same coordinates could be written in two different files. It is therefore an additional empirical regularity, rather than a necessary truth, that different pseudo-objects have different pseudo-locations, that is, that they are not, so to speak, “double-booked.”²⁰

The details of this argument raise a number of questions.²¹ But even if this particular line of thought doesn't serve as a definitive response to the skeptic, it illustrates a crucial difficulty the skeptic faces. Some relations between properties are necessary. To that extent, properties have what might be called a “modal configuration.” If F and G are properties whose modal configurations are different, one can't simply substitute reference to G for reference to F within a hypothesis, and assume that the structure of the original can be preserved. My argument about the role of location-properties in the RWH was meant to exploit that fact. The idea was that the modal configuration of location-properties includes their being uniquely instantiated, that is, that necessarily, if x is located at L, then nothing else can be located at L. No other genuine property P is such that, necessarily, if x is P, then nothing else is P. Hence, a version of the ISH which adverts to other properties in place of location-properties must add some empirical “exclusion principle” for these pseudo-locations to explain why no two items share the same pseudo-location. The upshot would be that substitution of reference to pseudo-locations rather than genuine locations in the ISH doesn't permit preservation of the explanatory structure of the RWH. More generally, the ISH will fail to match the RWH in explanatory adequacy unless the properties invoked by the ISH have exactly the right modal configurations. I see no reason to believe that this requirement can be met, and substantial reasons to doubt it.²² Considerations of explanatory adequacy will, then, favor the RWH over the ISH.²³

Here is where we stand. Skeptical hypotheses usually discussed, such as the MSH, are underdeveloped, explanatorily impoverished, and thus defective. The question then arises whether there is a superior skeptical hypothesis which could compete on equal terms with the RWH. My answer is no.

III

The ground gained so far is significant, but not unlimited. We have a refutation, or the outline of a refutation, of a certain kind of skepticism: domestic skepticism about the external world, motivated by the underdetermination principle. I have assumed that, since the skepticism under consideration is domestic skepticism, a reply may freely appeal to whatever epistemic principles we ordinarily accept, including principles of IBE. But the status of IBE itself is controversial, and some will doubt that an appeal to explanatory considerations has much force against the skeptic.

A principal criticism directed against IBE is that epistemic justification would be worthless unless beliefs that are so justified are thereby more likely, or have a greater tendency, to be true. Now, suppose H_1 provides appropriately simpler, hence better, explanations of a given domain than H_2 does, so that IBE underwrites the choice of H_1 over H_2 . This choice seems to presume that the simplicity of an explanation is a guide to truth, that is, that the world is likely to be simple in the ways we appreciate. But, the criticism goes, such a presumption is nugatory – it is a kind of wishful thinking that the world is nice and neat. Since simplicity and truth seem to be independent, explanatory success gives us at best pragmatic, rather than epistemic, reasons to favor one hypothesis over another. The upshot for skepticism is that, if the hypothesis that you are the victim of massive sensory deception is more complex than the real world hypothesis, the former may suffer some pragmatic disadvantage (it may be more cumbersome, inelegant, or the like). Nevertheless, you don't have genuinely epistemic justification for rejecting that hypothesis, and the deceiver argument can still go through.

The issues raised by this kind of argument are various and difficult, and I can't hope to do them justice on this occasion. But I would like at least to register the opinion that, in the end, IBE is no more suspect than other kinds of inductive confirmation. IBE licenses the acceptance of hypotheses that aren't entailed by the evidence that supports them, as any sort of inductive confirmation does. For example, consider an application of induction by enumeration. Observations that the roses in the yard have all been red in the past is evidence for the conclusion that the roses in the yard will be red next season. But the proposition that all the roses in the yard have been red in the past doesn't entail that all the roses in the yard will be red next season. IBE and enumerative induction have still more in common. The criticism of IBE just considered was that acceptance of a hypothesis because of its simplicity presupposes that simplicity is a guide to truth, and that such a presupposition amounts to a kind of wishful thinking that the world is nice and neat. But one might say with equal justice that enumerative induction presupposes that the observed is a guide to the truth about the unobserved, and that this presupposition, too, is a kind of wishful thinking that the world is nice and neat. To this extent, IBE appears to be on a par with other forms of inductive confirmation.²⁴

Given the assumption that knowledge requires justification, there are now three positions in play: skepticism about the external world, skepticism about IBE, and skepticism about induction in general. In my view, the relations among them are as follows. If you aren't a skeptic about IBE, you can reject skepticism about the external world. If, however, you are skeptical about IBE for the reasons given above, you will also be a skeptic about induction of all kinds (and presumably about the external world, too).

So far, I haven't offered any reply to skepticism about induction, nor will I attempt to do so. Instead, I'll make some comments that are meant to illuminate, rather than settle, some of the issues now before us.²⁵

(1) One can take a sunny view of the situation, according to which there can be no rationally compelling challenge to our basic epistemic principles. A critic might reject our principles, or be committed to principles that conflict with ours. But the mere fact that someone else disagrees with us hardly shows that we must be wrong. If anything, we should stand by our own principles, and regard the critic's stance as mistaken. Assuming then that our epistemic principles include or imply the canons of IBE and induction in general, their legitimacy isn't open to serious question. To use the terminology introduced above, skepticism about IBE is exotic and, consequently, may be ignored. The same goes for skepticism about induction more broadly.²⁶

(2) This response may strike some as unsatisfactory. For one thing, it may be disappointing that we are unable to establish the correctness of our epistemic principles vis-à-vis others, from some kind of "neutral" standpoint. In addition, the line of criticism directed against IBE and induction at the beginning of this section may seem to have some weight, and just setting it aside may appear superficial or obtuse. But exactly what is at issue here is somewhat cloudy. One sharpening of the concern about the connection between justification and truth is what I have termed the *problem of misleading evidence* (see Vogel, 2004, 2007). Suppose your evidence E inductively confirms hypothesis H, so that you're justified in believing H. Presumably, you're justified in believing propositions logically weaker than H, including $(\tilde{E} \ \& \ \tilde{H})$.²⁷ If $(E \ \& \ H)$, then H is false despite your having good evidence for H. We could say that, in such a case, your evidence is misleading with respect to H. Now, we're supposing that you're justified in believing $(\tilde{E} \ \& \ \tilde{H})$, that is, that your evidence *isn't* misleading with respect to H. That is a proposition about the world, so you're not justified in believing it unless you have evidence that supports it. It may appear, though, that the only pertinent evidence you have is E. That looks like trouble, because it is hard to see how E could be evidence that E itself isn't misleading with respect to H. But if you have no evidence at all for $(\tilde{E} \ \& \ \tilde{H})$, it seems that you're not justified in believing $(\tilde{E} \ \& \ \tilde{H})$. It follows that you're not justified in believing H, and the result is a thoroughgoing skepticism about induction. As far as I'm concerned, whether the problem of misleading evidence is a serious threat to inductive knowledge, and, if so, how to respond to it, remain open questions.

(3) In his contribution to this volume, Richard provides a searching and distinctive treatment of motivations for a general skepticism about (non-deductive) "inferential justification." According to Richard, someone who rejects externalist accounts of justification must accept that

When one's justification for believing P involves *inference* from E [one's evidence for P], a constituent of that inferential justification is a justified belief that E does indeed make probable P ... We could, of course, *infer* that the probability connection holds from some other proposition we believe, F, but then we would still need reason to think that F makes likely that E makes likely P. Eventually, however, inferential internalists will need to find some proposition of the form E makes likely P that we can accept without inference.

Substantive propositions of that sort look to be very hard to come by.²⁸ In that event, skepticism about induction immediately follows. Now, I'm not confident that I grasp exactly what Richard means by "makes probable," but structurally this challenge resembles the problem of misleading evidence. The former makes having evidence for "E makes probable H" a condition for justified belief that H; the latter makes having evidence for "E isn't misleading with respect to H" a condition for justified belief that H. The two requirements are also similar in spirit, or so it seems to me.²⁹ Thus, I'm inclined to think that there will be a satisfactory response to Richard's challenge if and only if there is a satisfactory response to the problem of misleading evidence – assuming that a response is necessary at all.³⁰

This reservation reflects uncertainty as to the status of the demand that, if S's believing E justifies S's believing H, then S must have evidence for the proposition that E makes probable H (hereafter, "Fumerton's requirement," or FR). I take Richard's view to be that, given the falsity of externalism, FR holds. This means, at a minimum, that FR is a condition on what *we take to be* inferentially justified belief. FR is an epistemic principle we accept. Richard maintains that FR leads to skepticism, in the absence of a priori, non-inferential justification for judgments of the form *E makes probable H*. Suppose all that is so. The result is *domestic* inductive skepticism, insofar as such skepticism follows from epistemic principles to which we are committed.

Now, there is at least some reason for denying that FR is a condition for justification, as we understand it. Consider the total body of evidence (TE) available to a subject S. When S has inductive justification for a proposition H, TE confirms but doesn't entail H.³¹ According to FR, S still needs reason to think that TE makes probable H. And, since TE doesn't entail H, S presumably needs some additional *evidence* that the making-probable relation holds between TE and H. But S can have no such evidence, since, by assumption, TE is S's total evidence. FR, then, seems to be unsatisfiable in principle. So, if we have a *coherent* conception of justification, it can't include FR. That result, in turn, favors the view that skepticism stemming from FR is exotic, not domestic.³²

But maybe justification as we conceive it *is* subject to FR, and it undoes itself in the way just described. Our situation with respect to epistemic justification would then be something like the one we face with respect to the concept of a set, in the face of Russell's Paradox, or with respect to the concept of truth, in the face of the semantic paradoxes. After the bad news is in, we can try to minimize our disappointment by seeking out a successor concept that does a good part of the work done by its untenable predecessor. The question in the present case would then be whether our concept of justification, shorn of FR, is enough like the original for us to care about it as before. If so, the encounter with the sort of skepticism Richard envisions might end in partial concession rather than full capitulation.

IV

In this essay, I have argued that some kinds of skepticism need to be refuted, and can be. Other kinds of skepticism can't be refuted, but needn't be. I hope, and want to suggest, that the division is exhaustive.

Acknowledgments

I would like to thank Philip Bricker, Richard Feldman, Michael Glanzberg, Ned Hall, Mark Johnston, Tom Kelly, Tim Maudlin, Richard Moran, Bernard Nickel, Ted Sider, and Steve Yablo for their kind help. I have been very fortunate to work with Richard Fumerton, Ernest Sosa, and Matthias Steup on this project. Finally, I'm deeply grateful to Neil Harpster for his extraordinary efforts on the side of the angels.

Notes

- 1 Generically, skepticism is the denial that we have some kind of knowledge (or, maybe better, the denial that we have some kind of knowledge one might sensibly suppose we have). The rejected body of knowledge might be identified by subject matter (skepticism about the external world, skepticism about the past) or by mode of acquisition (perception, induction). When I talk about "skepticism" as such, I mean skepticism about the external world. In addition, I assume here that knowledge requires justification, so that the denial that we have justification for a body of beliefs implies the denial that those beliefs count as knowledge. For more on these matters, see Richard's helpful discussion, and section III below.
- 2 What follows is an updated version of Descartes's "Evil Genius" argument in the First Meditation.
- 3 For simplicity, I'm ignoring the possibility that the toaster shut off because it burned out *and* also because the fuse blew. I'm proceeding as though the two hypotheses are *competitors*, in the sense that the truth of one is logically incompatible with the truth of the other.
- 4 For discussion of these issues, see Vogel (2007).
- 5 Richard considers the investigation of various conditional claims of the form *If such and such principles hold, skepticism follows*. His characterization of the way I understand the situation is helpful, except for two things. First, I don't think I've misidentified the epistemic principles to which we're committed, being "stingy" about excluding "*sui generis* epistemic principles of perception" (see Vogel, 1997). Second, I'm not sure how Richard distinguishes "conceptual requirements for knowledge and justified belief" from "the legitimacy of certain basic sorts of inference." Thus, I don't see that there are two different ways of drawing a distinction between domestic and exotic skepticism.
- 6 For extended, critical discussion, see Williams (1992).
- 7 For extended, critical discussion, see Greco (2000).
- 8 For my evaluation of these views, see Vogel (1993, 1997).
- 9 Whether there is such a body of facts to be explained is disputed by Williams (1992) and Byrne (2004). For a response, see Vogel (1997).
- 10 Other authors have advanced different replies to skepticism that, in one way or another, appeal to IBE. Two recent contributions are those due to BonJour (1998) and Peacocke (2004).
- 11 For different views about what an explanation is, see Pitt (1988).
- 12 My discussion blurs some important distinctions. At least some explanations are explicitly causal, where the *explanans* is a cause of the *explanandum*. I assume that at least some explanations involve the subsumption of the *explanandum* under some law or regularity, either explicitly or because causation itself requires nomological support. In the text, I make no attempt to sort this out.
- 13 The terms of the causal relation are usually taken to be events, not things. The scrupulous reader may make the necessary adjustments to the text.

- 14 Otherwise, science would be even harder than it is. For example, proponents of steady-state cosmology would have been able to explain the presence of background radiation as well as proponents of the Big Bang. Then, too, the job of defense attorneys would be far easier than it actually is.
- 15 According to the view known as “structural realism,” we can know that things in the world have at least some relational properties, but we are unable to know what intrinsic properties they have. A structural realist might then deny that we have reason to accept the RWH rather than the ISH. I take it Richard is sympathetic to such a position. But if shape is an intrinsic property of bodies, then a structural realist of this stripe has to say that we don’t know what shapes things have, which is a more skeptical outcome than I would like to see. Matters here are complicated, however, involving among other things questions about the identity conditions for properties that I cannot discuss in any detail.
- 16 The difference in purely spatial features of our environment doesn’t exhaust the disagreement between the ISH and the RWH. But I think that if skepticism about the spatial configuration of the external world can be refuted, dealing with whatever else remains of external world skepticism would be relatively easy.
- 17 The example is from Lewis (1980).
- 18 By “distinct” I mean at least that the objects have no parts in common at the time in question. So, if you think that a statue and the mass of clay from which it is fashioned aren’t identical, the statue and mass of clay still wouldn’t count as distinct, in the way that I’m employing the term. If your favored ontology allows for object-stages, the thesis is meant to imply that there can’t be two distinct object-stages at the same place at the same time.
- 19 On the relation between simplicity and necessary truth see Glymour (1984), and below (compare Sober, 1988).
- 20 A bit more precisely: an account according to which X has the property of being at a location L will explain more economically why X behaves as though it has that property than an account according to which X really doesn’t have the property of being at location L, but rather has a pseudo-location which corresponds to L.
- 21 Mark Johnston and Philip Bricker have cautioned that there may be counterexamples to the claim that the co-location of distinct objects is impossible. Ned Hall and Tim Maudlin have suggested that the propensity of bodies to exclude one another is a matter of the dynamics governing those bodies, and thus not necessary.
- 22 Here is the sort of thing I have in mind. I am inclined to think that spatial properties are such that, necessarily, they conform to some kind of geometry, and their so doing enters into the modal configurations of those properties. A relatively clear – although somewhat out of the way – case in point may be the following. Consider the property of being a perfect solid. In Euclidean space, there are exactly five perfect solids, so the determinable *perfect solid* has precisely five determinates. Suppose the RWH invokes the fact that something is a perfect solid in giving an explanation (say, of why the chances that each side of an object will come up on a toss are equal). To maintain structural parity with the RWH, the ISH would have to invoke some other property (the property of being a “pseudo-perfect-solid”) with exactly five determinates, while avoiding other difficulties. I see no reason to think that this condition can be met. The sort of difficulty I am raising here is analogous to that facing the possibility of spectrum inversion: color-properties bear complex relations to one another, such that a simple substitution of one for another in an account of color experience may not be feasible. See Hardin (1997).
- 23 There may be another way of arguing for the same result, which moves in something like the opposite direction from the foregoing. The ISH isn’t merely a competitor to the RWH, but is something like a simulacrum of it. As the ISH more closely resembles the RWH in structure, the question of what makes the ISH *different* from the RWH becomes more pressing. The problem, loosely speaking, is that the ISH needs to specify how things that behave as though they are F still *aren’t* F, without taking on unwanted encumbrances.

- 24 There are other notable objections to IBE, including those due to Van Fraassen. For a survey and discussion of these, see Ladyman et al. (1997). I should note that Van Fraassen would reject the view that IBE is no worse off than any other kind of ampliative inference, and that Richard agrees with Van Fraassen at least that far.
- 25 One important difference is that the deceiver argument works by invoking the underdetermination principle. Richard examines skepticism about induction (and the external world) motivated by what he calls the principle of inferential justification. For a discussion of these issues, see Vogel (2007), and below.
- 26 This position is similar to the one adopted by P.F. Strawson (1952), although it doesn't include Strawson's view that claims about inductive support are analytic. Strawson undertakes to show in some detail how various motivations for inductive skepticism covertly distort or diverge from our ordinary understanding of what is required for epistemic justification. Those efforts bolster the judgment that skepticism about induction is invariably exotic in character.
- 27 That is, it is not the case that E is true while H is false. This is a more precise rendering of what it is for the world to be "nice and neat."
- 28 Richard also stresses that the justification for the underpinning propositions must be a priori, which if anything makes it less plausible that there are such propositions.
- 29 See Vogel (2007) for discussion of what I take to be the likeness of the two. But there are differences, as I note there.
- 30 But some caveats are called for. First, adopting some kind of holism might block the threat Richard envisions, although it is somewhat harder to see how holism could help with the problem of misleading evidence. Second, the link between justification for H and justification for (\tilde{E} & \tilde{H}) is secured by the eminently plausible principle that justification is closed under known logical implication. What underwrites the connection between justification for H and justification for (E makes probable H) is somewhat harder to see, and is perhaps more tenuous.
- 31 I take it Richard would not protest at this point, in light of his reply to a criticism by Michael Huemer.
- 32 *One* way to read FR is that it covertly appeals to a deductivist view of confirmation, as follows. Whenever E is inductive evidence for H, FR finds E insufficient, and demands further evidence (for the proposition that E makes probable H). So, the import of FR is that inductive evidence is never sufficient for justification; only deductive evidence would really do. Richard agrees that the standard of justification for knowledge should not rise that high, although I take it he would reject the way I have characterized FR. Points related to this one, and the one in the text, may also provide some refuge from the problem of misleading evidence.

References

- BonJour, L. (1998) *In Defense of Pure Reason*. Cambridge: Cambridge University Press.
- Byrne, A. (2004) How hard are the skeptical paradoxes? *Noûs* 38: 299–325.
- Glymour, C. (1984) Explanation and realism. In *Scientific Realism*, ed. J. Leplin. Berkeley: University of California Press.
- Greco, J. (2000) *Putting Skeptics in Their Place*. Cambridge: Cambridge University Press.
- Hardin, C.L. (1997) Reinverting the spectrum. In *Readings on Color, Volume I*, ed. A. Byrne and Hilbert. Cambridge, MA: MIT Press.
- Ladyman, J., Douven, I., Horsten, L. and van Fraassen, B. (1997) A defence of Van Fraassen's critique of abductive inference: reply to Psillos. *Philosophical Quarterly* 47: 305–321.
- Lewis, D. (1980) Mad pain and Martian pain. In *Readings in Philosophy of Psychology, Volume I*, ed. N. Block. Cambridge, MA: Harvard University Press.

- Peacocke, C. (2004) Explaining perceptual entitlement. In *The 'Challenge' of Externalism*, ed. R. Schantz. Berlin: de Gruyter.
- Pitt, J., ed. (1988) *Theories of Explanation*. Oxford: Oxford University Press.
- Pryor, J. (2000) The skeptic and the dogmatist. *Noûs* 34: 517–549.
- Sober, E. (1988) Let's razor Occam's razor. In *Reconstructing the Past: Parsimony, Evolution, and Inference*. Cambridge, MA: MIT Press.
- Strawson, P.F. (1952) *Introduction to Logical Theory*. London: Methuen.
- Vogel, J. (1990) Cartesian skepticism and inference to the best explanation. *Journal of Philosophy* 87: 658–666.
- Vogel, J. (1993) Dismissing skeptical possibilities. *Philosophical Studies* 70: 235–250.
- Vogel, J. (1997) Skepticism and foundationalism: a reply to Michael Williams. *Journal of Philosophical Research* 11–28.
- Vogel, J. (2004) Subjunctivitis. *Philosophical Issues* 14: 426–455.
- Vogel, J. (2007) Skeptical arguments. *Philosophical Studies* 134: 73–88.
- Williams, M. (1992) *Unnatural Doubts*. Oxford: Blackwell.
- Wittgenstein, L. (1972) *On Certainty*. New York: Harper and Row.

The Challenge of Refuting Skepticism

Richard Fumerton

One of the keys to success in battle is to find and hold ground favorable to the engagement. The skeptic, it is often claimed, insists on unacceptably strong standards for knowledge or justified belief and proclaims victory when those standards are not met. In his paper Jonathan distinguishes what he calls *exotic* skepticism from what he takes to be its more interesting *domestic* versions. I join with Jonathan in an attempt to make clear what a skeptic might reasonably demand of a successful refutation of skepticism. As will be obvious, Jonathan and I agree on a great many issues, but may part company on just exactly how to understand epistemic probability and the need for access to that probability as a requirement of acquiring ideal justification.¹ I will conclude with a brief discussion of the prospects of avoiding skepticism within the framework of the domestic skepticism that Jonathan discusses.

Defining the Skeptical Challenge

When addressing a skeptical challenge one should probably first attempt to discover whether the skeptic is bent on undermining *knowledge*, or has the more ambitious goal of undermining *justified belief*. We can also distinguish *global* skepticism from its *local* varieties. The global skeptic with respect to *knowledge* claims that we don't know anything – we can't even know that we don't know anything (though we might be justified in believing that). The global skeptic with respect to *justified belief* claims that we have no justified beliefs – we are not even justified in believing that we have no justified beliefs. There have been relatively few global skeptics with respect to knowledge and fewer still with respect to justified belief. The global skeptic with respect to justification

must obviously deal with the charge of epistemic self-refutation. Any argument for global justification skepticism will have premises that cannot be justifiably asserted if the conclusion of the argument is true. It is not clear that that fact entitles one to ignore the arguments, but it puts the most extreme of skeptics in an awkward position to say the least. In any event, even rather radical skeptics tend to at least implicitly suppose that they can know truths about their subjective experience and recognize the distinction between legitimate and illegitimate reasoning.

Local skeptics are skeptics with respect to knowledge or justified belief about some particular *kind* of truth: for example truths about the external world, the past, the future, other minds, theoretical entities posited by physics, and so on. Like Jonathan, I will often take, as our paradigm of a local skepticism, skepticism with respect to truths about the external world. The skeptic with respect to knowledge of the external world claims that one cannot know (or sometimes with added emphasis, know *with certainty*) any truths about the physical world. Before investigating this claim, we should decide what conditions we need to meet in order to satisfy our skeptic that we do indeed possess the knowledge in dispute. Suppose that our skeptic insists that we know only when we possess justification that precludes any possibility of error. If we agree, the skeptic might seem to face relatively smooth sailing. The justification I have for believing that I am seated in front of a computer is precisely the same sort of justification that I would have had were there no computer there and were I instead suffering a massive and vivid hallucination (or were I a brain in a vat, a victim of demonic machination, an inhabitant of the *Matrix* world, etc.). My justification for believing that the computer exists, therefore, does not entail the existence of the computer.

The above argument is not uncontroversial. Williamson (2000), Brewer (1999), and others, for example, argue that we might be unable to *recognize* crucial and dramatic differences between epistemic situations in which we find ourselves. We might not be able to tell, for example, the difference between a situation in which we are directly confronted with the computer (a situation in which the computer itself is literally a *constituent* of our visual experience) and a situation in which we are merely hallucinating the computer. Our perceiving that there is a computer in front of us might *in fact* generate a kind of justification that guarantees the computer's existence and, trivially, provide us with evidence of precisely the sort Descartes sought that guarantees the computer's existence.² For those of us who find the skeptic's challenge interesting, however, it is difficult to take this idea seriously. While we may be unable to tell the difference between epistemic situations that are only marginally different, it still seems utterly puzzling as to how we could be so incompetent as to fail to notice the difference between two experiential situations that Williamson (and other direct realists) claim should get *radically* different philosophical analyses.

In any event, for the purposes of this debate I will simply presuppose (and Jonathan, I understand, will not dispute) that we can easily imagine skeptical scenarios that are epistemic counterparts to veridical perception, scenarios in which we would possess precisely the same kind and strength of justification for the same kind of beliefs about our physical environment that we would possess were our perceptions veridical. We will agree here, therefore, that we cannot avoid knowledge skepticism with respect to the physical world if we understand knowledge as requiring justification so strong that it eliminates the possibility of error.

Of course, the moral most contemporary epistemologists draw is that we should reject Cartesian standards for knowledge. Our only hope of avoiding skepticism with respect to knowledge of the external world (the past, the future, other minds, and so on) is to accept more relaxed standards for knowledge. While Jonathan and I may end up disagreeing on whether or not one can defend weaker requirements for knowledge, that disagreement is probably unimportant for this debate. I think we do agree that the most interesting epistemological question is whether we can find justification that makes likely for us the truth of what we believe. As a result I'm going to frame the skeptical challenge employing the concept of justification with the understanding that if one can know without possessing justification that guarantees the truth of what one believes, the discussion could just as easily take place employing the concept of knowledge.

Underdetermination

If we switch our focus from knowledge to justified belief then what should one be expected to show in order to defeat skepticism? Here, I agree completely with Jonathan that the issue is best understood in terms of underdetermination. He suggests two principles. While he states the principles in terms of knowledge, one can easily reformulate them to concern justification instead:

- UP1: If q is a competitor to p , then there is justification for one to believe p only if one can non-arbitrarily reject q .
- UP2: If q is a competitor to p , then there is justification for S to believe p only if p has more epistemic merit for S than q .

For these purposes, competitors to p include all contraries of p and any other proposition whose truth would render highly unlikely the truth of p .³

UP1 and UP2 state only a *necessary* condition for S 's possessing justification for believing P . A skeptic might want to remind us that for S to possess justification for believing P , there must be an epistemically non-arbitrary way of rejecting the *disjunction* of competitors to p . This is particularly important to remember in evaluating arguments to the best explanation. To have justification for accepting a given explanation h for some data e one must have more reason to accept h than the *disjunction* of competing explanations. When the skeptic comes up with competitors to the hypotheses of commonsense beliefs, therefore, one must be very careful not to assign these competitors *any* sort of significant possibility of being true. The disjunction of a large number of very unlikely hypotheses is still often very likely to be true.

The nice thing about framing the debate in terms of underdetermination is that we beg no question against the skeptic. We meet the skeptic on a level playing field. If we are to take seriously the epistemic challenge of skepticism, we must insist on no burdens of proof. The beliefs of common sense are not epistemically innocent until proven guilty, and skeptical scenarios are not epistemically guilty until proven innocent.

Choosing from among Competitors

How are we supposed to figure out whether a hypothesis of common sense satisfies the above requirements? Well, if we are foundationalists, we need both a meta-epistemological account of non-inferential justification and a normative epistemological account of what we are non-inferentially justified in believing.⁴ If beliefs about the external world end up in the class of beliefs that are non-inferentially justified, and the skeptic's competitors to those beliefs do not, then the skeptic loses. One takes the propositions of common sense and employs closure to reject the skeptic's competitors.⁵

Given my own views about non-inferential justification, propositions about the physical world (the past, the future, other minds) do *not* get included in the class of propositions for which we possess non-inferential justification. I can't argue here for my version of foundationalism. The view I defend (1995 and elsewhere) is very traditional and now very unpopular. In short, I hold that one is only non-inferentially justified in believing a proposition *p* when one is directly acquainted with the fact that *p* while one has the thought that *p* and one is acquainted with a correspondence between the thought and the fact. Because I think that one can possess in hallucinatory experience the *same* justification that one possesses in veridical experience for believing truths about one's physical environment, I don't think one can be directly acquainted with facts about physical objects.

Jonathan wants to remain neutral on many of the most fundamental questions concerning the *structure* of justification (the debate between foundationalists and coherentists), and even neutral on the question of what we should take to be the class of beliefs non-inferentially justified if we endorse foundationalism. But because he thinks that his approach to meeting the skeptical challenge through argument to the best explanation works from within quite different frameworks, I think he would not object to my presupposing for the purposes of this debate a foundationalism that takes as its base propositions restricted to the phenomenal (the world of appearance).

So how do we move beyond our foundational justification for believing some proposition *e* to justified belief in a proposition of common sense, say a proposition describing our immediate physical environment? The two most obvious suggestions are these:

- 1: For *S* to have justification for believing *p* on the basis of *e*, *e* must make epistemically probable *p* (where *e*'s entailing *p* can be viewed as the upper limit of *e*'s making *p* probable).
- 2: For *S* to have justification for believing *p* on the basis of *e*, *S* must be aware of the fact that (have justification for believing that) *e* makes probable *p*.

Elsewhere, I have called (2) *inferential internalism* (see Fumerton, 1995, chapter 3). Notice that the inferential internalist does *not* endorse the general principle that to be justified in believing *P* one must be justified in believing that one is justified in believing *P*. Nor does he endorse the principle that to be inferentially justified in believing *P* one must be justified in believing that one is inferentially justified in believing *P*. *Those* principles seem destined to generate vicious regress.⁶ The inferential internalist insists only that when one's justification for believing *P* involves *inference* from *E*, a *constituent* of that inferential justification is the justified belief

that E does indeed make probable P. One must see the connection between one's evidence and one's conclusion before one can justifiably draw the conclusion from one's evidence.

If we are foundationalists and we accept inferential internalism, then there isn't just one potentially vicious regress of justification we need to end with non-inferential justification. To be sure, in order to have inferential justification for believing P, we must be able to trace our justification for believing P back to some proposition for which we have non-inferential justification – a proposition that will serve as our premise. But we also need to cauterize a threatening regress concerning the justification we need for believing that our premise makes probable our conclusion. We could, of course, *infer* that the probability connection holds from some other proposition we believe, F, but then we would still need reason to think that F makes likely that E makes likely P. Eventually, however, inferential internalists will need to find some proposition of the form E makes likely P that we can justifiably believe without inference. And, indeed, the *key* to meeting the skeptical challenge for the inferential internalist centers on the ability to find non-inferential justification for accepting probability connections between our available evidence and the propositions of common sense.

Leaving aside for now the issue of how to understand epistemic probability, why should one embrace inferential internalism? Why shouldn't we view the inferential internalist's demands on inferential justification as exotic, as demands that simply *invite* skepticism? Well, consider the astrologer who predicts the future based on beliefs about the positions of planets relative to one another. It seems obvious that a *sufficient* condition for rejecting astrological inferences as rational is that the astrologer has no reason to believe that the positions of planets relative to one another makes *likely* any of his predictions. If a detective infers from someone's appearance that the person is guilty of a crime, one will surely demand evidence for supposing that appearances of this kind are correlated with guilt before one will concede the rationality of the conclusion. These commonplace examples and indefinitely many others like them surely indicate that we *do* embrace the inferential internalist's account of what is necessary for inferential justification.

Mike Huemer (2002) has objected plausibly to the above arguments for inferential internalism. He argues that the examples used to make initially attractive the principle are misleading in that they inappropriately characterize the evidence from which one infers the relevant conclusion. Even astrologers don't think that they can legitimately infer their predictions from propositions describing the positions of planets and the birth dates of people, *and from that information alone*. It should be a truism that much of the argument we actually give outside of a philosophical context is highly compressed, highly enthymematic. As we ordinarily use the term "evidence," we certainly do characterize litmus paper's turning red in a solution as evidence that the solution is acidic. The approach of very dark clouds is evidence of an approaching storm. A footprint on a beach is evidence that someone walked on the beach recently. But it is surely obvious upon reflection that one's evidence for believing that the solution is acidic, for example, is not the color of the litmus paper *by itself*. To legitimately draw the conclusion one would need an additional *premise*, most likely a premise describing a correlation between the color of litmus paper in a solution and the character of that solution.

Once one realizes that the reasoning discussed above is enthymematic, one is positioned to respond to that appearance of an argument for inferential internalism. It is necessary to have some justification for believing that there is a connection between positions of planets and the affairs of people, the approach of dark clouds and storms, footprints on a beach and the recent presence of people before drawing the respective conclusions, but *only* because propositions describing connections or correlations of the relevant sort are implicitly recognized as critical *premises*. Internalists and externalists alike share the foundationalist's insight that inferential justification is parasitic upon the justification we possess for believing the relevant premises of our arguments. If the astrologer is relying on an unstated, but critical, premise describing correlations between astronomical facts and the lives of people in reaching her conclusion, she will, of course, need justification for believing that *premise*. But that in no way suggests that when we have fully described *all* of the relevant premises from which a conclusion is drawn, we should require that the person who draws that conclusion have additional evidence for believing that the premises make probable the conclusion. The existence of the relevant connection between premises and conclusion is enough.

One can make just as strong a case for inferential internalism, however, by focusing on non-enthymematic reasoning. Consider the case of someone who infers P from E where E logically *entails* P. Is the inferential internalist right in maintaining that in order for S to believe justifiably P on the basis of E, S must be aware of the fact that (or at least have a justified belief that) E entails P (or alternatively, that the inference in question is legitimate)? The answer still seems to me obviously yes. We can easily imagine someone who is caused to believe P as a result of believing E where E does in fact entail P, but where the entailment is far too complicated for S to understand. Unless S sees that P follows from E, would we really allow that the inference in question generates a justified belief? Or to make my case a bit stronger, would we allow that the person who reaches the conclusion has philosophically relevant justification or ideal justification – the kind of justification one seeks when one searches for philosophical assurance.

Summarizing, I'm not at all sure that the inferential internalist is imposing unreasonably exotic requirements on justification, at least if the justification we seek is justification that provides assurance from the first-person perspective. I would allow that one might well acknowledge *derivative*, less demanding concepts of justification. Elsewhere (Fumerton, 2004), I have argued that one might allow that when one is caused by the fact that E to believe P, when the fact that E is the truth maker for the proposition that E and the proposition that E makes probable P (either alone or with other propositions in our background evidence), one has a *kind* of justification for believing P. It is not the kind of justification I think the *philosopher* seeks, and to determine if we have it, we still need a justified belief of the sort sought by the inferential internalist.

The Analysis of Epistemic Probability

Whether or not we adopt inferential internalism, we need an analysis of the probability connection that by itself or as the object of awareness is partially constitutive of inferential justification. There are at least two quite different approaches one might take to

analyzing the epistemic probability with which philosophers are concerned. One approach attempts to understand what it means to say of our evidence that it makes probable a conclusion in terms of epistemic evaluation of belief. Crudely, E makes probable P when one's justification for believing E gives one justification for believing P (were E the only evidence one possesses). A philosopher taking this approach must analyze epistemic concepts without appealing to the concept of epistemic probability.

On the other approach, one appeals to an allegedly prior and more fundamental understanding of the relation of making probable holding between propositions in order to explicate inferential justification. On this approach, one's inferential justification for believing P on the basis of E is constituted by either the existence of, or, according to the inferential internalist, our *awareness* of, a probability connection holding between propositions. Just as one proposition or conjunction of propositions can entail another, so also one proposition or conjunction of propositions can make probable another. Just as there are deductively valid arguments whose deductive validity is not analyzable by epistemic concepts, so also there are valid non-deductive argument forms whose legitimacy depends on relations between the *contents* of premises and conclusions.

In what follows, I'm going to assume without argument that this last approach is correct. It is simply hard for me to believe that the justificatory status of inferentially justified beliefs is not fundamentally derived from relationships between that which is believed. Put another way, it is surely a feature of the *arguments* whose premises and conclusions are believed that is key to understanding the justificatory status of the beliefs formed in the conclusions as a result of justified belief in the premises.

But what is the best way of understanding the relation of making probable that holds between certain propositions? This debate has a long history, one that predates, but in many ways foreshadows, the now more familiar contemporary internalist/externalist controversies in epistemology.⁷ Painting with a very broad stroke, one can attempt to analyze probability claims in epistemology on the well known model of relative frequency that is offered as a way of interpreting claims about the probability of an individual or event having a certain characteristic. On a very crude interpretation of the frequency theory, to say of something that it is probably G is always elliptical for a more complex relativized claim of probability. One must refer the individual *a* that is G to some reference class F, and the more perspicuous statement of the probability claim is one about the probability of *a*'s being G relative to its belonging to the class F. The truth conditions for the claim of relative probability are determined by the percentage of Fs that are G. The higher the percentage of Fs that are G the more likely it is that something is G relative to its being F.

One could borrow at least the spirit of the relative frequency interpretation of probability and apply it to relations between propositions in the following way. We could suggest that in claiming that P is probable relative to E we are simply asserting that E and P constitute a pair of propositions, which pair is a member of a certain class of proposition pairs such that, when the first member of the pair is true, usually the second is. Thus in saying that *a*'s being G is probable relative to its being F and most observed Fs being G, I could be construed as claiming that this pair of propositions is of the sort: most observed Xs are Y and this is X/this is Y, and most often it is the case that when the first member of such a pair is true, the second is. Similarly, if I claim that my seeming to remember eating this morning (E) makes it likely that I did eat this morning (P), I could be construed as asserting that the pair of propositions

E/P is of the form S seems to remember X/X, such that most often when the first member of the pair is true, the second is.

The above view obviously resembles, at least superficially, the reliabilist's attempt to understand justified belief in terms of reliably produced beliefs. And it encounters many of the same difficulties. Just as the relative frequency theory of probability must inevitably move beyond actual frequencies in defining probability, so both the above account of epistemic probability and the reliabilist will inevitably be forced to move beyond actual frequencies in order to define the relevant epistemic probability/reliability. Just as reliabilism must deal with the generality problem, so the above approach to understanding epistemic probability as a relation between propositions must deal with the problem of how to choose from among alternative ways of characterizing the class of propositions pairs to which a given pair belongs. Just as many reliabilists are troubled by the implications of their view for what to say about worlds in which demons consistently deceive epistemically faultless believers, so a frequency theory of epistemic probability must deal with similar alleged counterintuitive consequences about what is evidence for what in demon worlds. Lastly, and most importantly for this debate, both reliabilism and the frequency theory of epistemic probability will be anathema to the inferential internalist who is convinced that one needs non-inferential *access* to probability connections in order to gain philosophically satisfying inferential justification. The inferential internalist who is a foundationalist will need to end a potential regress when it comes to gaining access to probabilistic connections. If one's model for foundational knowledge is something like knowledge of truths made true by facts with which one is directly presented, there seems no hope that one will get *that* kind of access to either the reliability of a belief-forming process or a probability relation (understood in terms of frequency) holding between propositions.

One of the historically most interesting alternatives to the frequency interpretation of epistemic probability is a view developed some 80 years ago by Keynes (1921). Keynes wanted to model epistemic probability on entailment. He held that just as one can be directly aware of the entailment holding between two propositions, so one can also be directly aware of a relation of making probable holding between two propositions. There are, of course, obvious differences between entailment and making probable. From the fact that P entails Q it follows that the conjunction of P with any other proposition entails Q. From the fact that P makes probable Q, it doesn't follow that P together with anything else makes probable Q. But for all that, we could still take making probable to be an a priori internal relation holding between propositions (where an internal relation is one that necessarily holds given the existence and non-relational character of its relata). P and Q being what they are, it cannot fail to be the case that P makes probable Q. (It might also be true that P, R, and Q being what they are it cannot fail to be the case that (P and R) makes probable not-Q.)

Against the Keynesian, one might argue that it is patently absurd to suppose that making probable is an internal relation holding between propositions. Such a view yields the absurd consequence that claims about evidential connections are necessary truths knowable a priori. If anything is obvious it is that the discovery of evidential connections is a matter for *empirical* research. While the objection might seem initially forceful, one must remember the point we conceded in considering Huemer's objections to inferential internalism. There is certainly no necessary evidential connection between litmus paper's turning red in a solution and the solution's being acidic, between

dark clouds and storms, between footprints on a beach and the prior presence of people. But then on reflection, we decided that it is misleading to characterize the litmus paper, dark clouds and footprints as the evidence from which we infer the respective conclusions. What we call evidence in ordinary parlance is just a *piece* of the very elaborate fabric of background information against which we draw our conclusions. So we shouldn't *expect* to find Keynesian probabilistic connections holding between, for example, the proposition that the litmus paper turned red and the proposition that the solution is acidic.

Where *should* we look for a plausible example of Keynes's relation of making probable? The obvious, though perhaps not all that helpful, answer is that we should look for it wherever we have what we take to be legitimate, *non-enthymematic* and *non-deductive* reasoning. The trouble, of course, is that philosophers don't agree with each other about what constitutes legitimate though deductively invalid reasoning. One might look at the relationship between the premises and conclusion of an enumerative inductive argument. Less plausibly, perhaps, one might think about the connection between the proposition that I seem to remember having an experience and the proposition that I had the experience. Still more problematically, we might suggest that my seeming to see something red and round necessarily makes probable that there is something red and round.

The view that there is an internal relation of making probable that holds between propositions is just what the inferential internalist *desperately* needs in order to avoid vicious regress. As we saw, if inferential internalism and foundationalism are true, then unless we are to embrace a fairly radical skepticism, we must find some proposition of the form E makes probable P that we can justifiably believe without inference. Since most foundationalists will concede that there are at least some propositions of the form E *entails* P that one can know without inference, the closer we can make our analysis of making probable resemble our analysis of entailment, the more plausible will be the claim that we can know without inference propositions of the form E makes probable P.⁸ As I implied earlier, I think that *if* there is a solution to skepticism it involves the ability to know a priori epistemic probability claims. But that is a very big *if* indeed, and the reason I think skepticism looms so ominously on the horizon is that it is difficult to convince oneself that one is acquainted with the relevant probability connections.

Another Sort of Domestic Skepticism

Suppose one cannot convince oneself either that there are relations of epistemic probability holding between propositions or that we have access to such relations. Is there further discussion we might have with the skeptic? Jonathan and I agree that skepticism about the knowledge that Descartes sought is not worth fighting. I'm not sure whether or not we agree that skepticism that presupposes inferential internalism is interesting. In any event, *these* issues concern the potentially exotic nature of *conceptual* requirements for knowledge and justified belief. As I understand Jonathan, however, there is yet another way of distinguishing between exotic and domestic skepticism. The domestic skeptic is prepared to *give* us the legitimacy of certain basic sorts of inference and will argue that even with these argument forms at our disposal we will be unable to generate

the conclusions of common sense from the available evidence. Now the inferential internalist discussed above will allow the use of an argument form only if you can know that the premises make probable the conclusion. You can have inferences from memory, for example, only if you know that your seeming to remember that you had some experience makes likely that you had the experience. You can legitimately employ inferences of enumerative induction only insofar as you can discover the probability connection between premises and conclusions of inductive inference. And the same would be true of the reasoning to the best explanation that Jonathan sees as our best hope of getting justification for beliefs about our physical environment.⁹ Before we can even assess the plausibility of the claim that the premises of arguments to the best explanation make probable their conclusions, we would need, of course, a very clear understanding of the *structure* of such arguments.

Whether or not there exists a rich enough array of legitimate and accessible non-deductive reasoning to generate conclusions of common sense, we certainly *can* still ask *conditional* questions of the following form: *assuming that* this or that form of reasoning is legitimate, can we non-arbitrarily choose the hypotheses of common sense over various skeptical hypotheses? Of course, if one includes among the inference forms described in the antecedents of the conditionals Chisholm-style epistemic principles whose formulations are tailor-made to avoid skepticism, the project might seem to lose interest. If, for example, it is conceded that it is epistemically permissible to infer that there is something red and square before one when it appears to one as if there is something red and square, then it will come as no surprise that sense perception will generate many of the conclusions of common sense. If one tries to get reasons for believing in the physical world employing only inferences of memory and enumerative induction, the project will be uphill. If one adds to these sorts of available arguments, reasoning to the best explanation, perhaps one might stand a better chance.

Again, I would emphasize that a reasonable skeptic should have no objection in principle to carrying on discussion concerning the plausibility of relevant *conditionals* asserting what one would be justified in believing were certain inferences legitimate. That same skeptic has every right to insist, however, that one can and ultimately *must* ask questions concerning the legitimacy of the reasoning referred to in the antecedents of the conditionals.

We can also make a distinction between domestic and exotic skeptics based on which principles of reasoning the skeptic is willing to presuppose as unconditionally true. We will, of course, still need some *principled* way of figuring out which presuppositions define the domestic skeptic. Jonathan obviously wants reasoning to the best explanation in the mix, and would probably appreciate an epistemic principle licensing conclusions about the past based on memory. He *wouldn't* insist that the domestic skeptic grant us *sui generis* epistemic principles of perception. One philosopher's domestic species of skeptic, however, is another philosopher's exotic skeptic. Jonathan would be characterized by many (not by me) as being unreasonably parsimonious in the epistemic paths he allows us to travel in an attempt to leave our foundations. One *could* try to single out the relevant presuppositions that define reasonable domestic skeptics by appealing to psychological facts about what inferences people actually employ, but here one must surely worry about the possibility that we find ourselves immersed in an irrational society that employs without hesitation fallacious reasoning.

Reasoning to the Best Explanation

I haven't addressed the question of whether one could defeat skepticism with the array of argument types Jonathan wants the skeptic to allow us. As Hume (1888, p. 212) so eloquently argued, we don't have a hope of getting justified belief about the external world if all we can rely on is memory and enumerative induction (where the premises of our inductive arguments are restricted to descriptions of correlations between experiences). I have argued elsewhere (1980) that there may be no legitimate form of reasoning to the best explanation understood as an alternative to inductive reasoning. I have further argued (1992) that even if we suppose that reasoning to the best explanation is an independent and legitimate argument form, it is far from clear that we can use it to arrive at the truths of common sense. In particular, it is hard to see how Berkeley's hypothesis that God directly produces in us sensations that come and go in certain predictable ways is inferior to the hypothesis that it is material objects playing that causal role. Berkeley himself claimed that his view was simpler. While he posited just two kinds of things – minds and ideas/sensations – the materialist posits three – minds, ideas/sensations, and material things.¹⁰

So again, *if* we give ourselves access to past experience through memory, enumerative induction as a way of projecting past correlations found between experiences into the future, and the prima facie plausibility of the claim that there are causal explanations for the occurrence of contingent phenomena, and, perhaps, even that we know somehow that simpler, more comprehensive explanations are true more often than complex, less comprehensive explanations, what are the prospects of successfully arguing for the commonsense view of the world over skeptical alternatives? Leaving aside the threat of Berkeley's idealism, a great deal depends on what we take propositions about the physical world to assert. If, for example, we embrace some form of representative realism and view common sense as vindicated only if there exist objects that are in some sense accurately *represented* by sense experience, then in addition to the competitors described by exotic skeptical scenarios, one must worry about the now familiar, but at one time quite disconcerting, fact that science *itself* seems to indicate that the physical world is quite different from anything we took it to be.¹¹ The solid table is mostly space. Its straight edges, when seen under high magnification, are decidedly crooked. The color that we think is *on* the object disappears under that same magnification. Although we are now used to these discoveries, one might well conclude that science has already undermined the common sense of *representative* realism.

On the other extreme, we might adopt Hylas's last suggestion (in the second of the *Dialogues*) before capitulating completely to Philonous's idealism. Hylas wondered why we couldn't think of matter as that, *whatever* it is, that is causally responsible for sensations coming and going in the familiar ways they do. If encouraged to pursue this idea, he would probably have been willing to offer further *negative*, and also very general *causal*, characterizations of these material things – they are not minds, there are many discrete objects, each with its own causal powers and so on. But on the view, our concept of a physical object involves no characterization of its *intrinsic, non-relational* properties. The world of external things is exhausted, both conceptually and epistemically, by their causal powers. On *this* conception of physical objects,

the prospects of avoiding skepticism are far brighter, but only because we require so little of the causes of experience in order for the beliefs of common sense to be vindicated.¹²

Whatever answer we give to these conditional questions, however, I want to stress one last time that we are only postponing the skeptic's final offensive. We will be asked about the truth of the antecedents of these conditionals, and we had better be prepared to indicate what we take their truth makers to be and how we think we can get access to them.

Notes

- 1 The fact that I refer to this justification as ideal might already be a warning that Jonathan (and others) may characterize its requirements as exotic.
- 2 On Williamson's view evidence is identified with knowledge. A perception of a computer provides a *context*, he thinks, in which one can know that the computer exists. That the computer exists is now part of our evidence and it trivially entails its own truth.
- 3 So, for example, the hypothesis that I am dreaming right now is a competitor to the proposition that there is a computer before me insofar as the truth of that hypothesis (together with the evidence in my possession) makes very unlikely the truth of the proposition that the computer exists. Alternatively, one might restrict competitors to contraries and include the negation of the proposition under skeptical attack as a conjunct of the description of the skeptical scenario.
- 4 If we are *not* foundationalists, we probably won't take the problem of skepticism seriously. The skeptic knows full well that skeptical scenarios are not going to cohere well with what we believe. I know from conversation with him that Jonathan isn't so sure that the skeptical problems disappear from within the framework of a coherence theory.
- 5 The closure principles states that if one is justified in believing P and knows that P entails Q then one has justification for believing Q. The principle of closure has been called into question, but the bottom line is that if a philosopher advances a view that forces us to reject closure, that should be taken as a *reductio* of that philosopher's view.
- 6 Again, for a discussion of reasons to reject global access requirements, weak and strong, for justification see Fumerton (1995, chapters 3 and 4).
- 7 One of the most interesting debates that has clear implications for the internalism/externalism controversy can be found in Keynes (1921) and Russell (1948, part V).
- 8 Notice that the Keynesian approach to understanding probability does not require that for one to be inferentially justified in believing P on the basis of E one must know or be able to formulate general principles of probability. One might be able to see the connection between particular propositions without seeing how to generalize. An analogous point holds of entailment. One can see that P entails Q without being able to see that entailment as an instance of *modus ponens*, *modus tollens* or any other general *kind* of entailment.
- 9 Vogel's (1990) defense of reasoning to the best explanation as the way in which to respond to the skeptic is by far the most careful and well developed in the literature.
- 10 Of course, a great deal depends on how one counts kinds. The mind of God, while a mind, isn't a whole lot like our minds.
- 11 See Eddington (1929) and Russell (1972, lecture IV) for further discussion of the idea that science undermines common sense.
- 12 For a defense of the claim that we do understand claims about the physical world this way, see Fumerton (1985, 2002, chapter 6).

References

- Brewer, B. (1999) *Perception and Reason*. Oxford: Oxford University Press.
- Chisholm, R. (1977) *Theory of Knowledge*, 2nd edn. Englewood Cliffs, NJ: Prentice-Hall.
- Eddington, A.S. (1929) *The Nature of the Physical World*. New York: Cambridge University Press.
- Feldman, R. and Conee, E. (1998) The generality problem for reliabilism. *Philosophical Studies* 89: 1–29.
- Fumerton, R. (1980) Induction and reasoning to the best explanation. *Philosophy of Science* 47: 589–600.
- Fumerton, R. (1985) *Metaphysical and Epistemological Problems of Perception*. Lincoln: University of Nebraska Press.
- Fumerton, R. (1992) Skepticism and reasoning to the best explanation. *Philosophical Issues* 2: 149–169.
- Fumerton, R. (1995) *Metaepistemology and Skepticism*. Lanham, MD: Rowman & Littlefield.
- Fumerton, R. (2002) *Realism and the Correspondence Theory of Truth*. Lanham, MD: Rowman & Littlefield.
- Fumerton, R. (2004) Achieving epistemic ascent. In *Sosa and His Critics*, ed. J. Greco. Oxford: Blackwell.
- Goldman, A. (1979) What is justified belief? In *Justification and Knowledge*, ed. G. Pappas. Dordrecht: Reidel.
- Goldman, A. (1986) *Epistemology and Cognition*. Cambridge, MA: Harvard University Press.
- Goldman, A. (1988) Strong and weak justification. *Philosophical Perspectives* 2: 51–69.
- Huemer, M. (2002) Fumerton's principle of inferential justification. *Journal of Philosophical Research* 27: 329–340.
- Hume, D. (1978) *A Treatise of Human Nature* (1888), ed. L.A. Selby-Bigge. Oxford: Oxford University Press.
- Keynes, J. (1921) *A Treatise on Probability*. London: Macmillan.
- Russell, B. (1948) *Human Knowledge: Its Scope and Limits*. New York: Simon and Schuster.
- Russell, B. (1972) *Our Knowledge of the External World*. London: Allen & Unwin.
- Vogel, J. (1990) Cartesian skepticism and inference to the best explanation. *Journal of Philosophy* 87: 658–666.
- Williamson, T. (2000) *Knowledge and Its Limits*. Oxford: Oxford University Press.
- Yourgrau, P. (1983) Knowledge and relevant alternatives. *Synthese* 55 (2): 175–190. Special Issue: *Justification and Empirical Knowledge*, ed. R. Audi.