

Skepticism, Sensitivity, and Closure, or Why the Closure Principle is Irrelevant to External World Skepticism

ADAM LEITE
*Department of Philosophy,
Indiana University, Bloomington*

Is there a plausible argument for external world skepticism? Robert Nozick's well-known discussion focuses upon arguments which utilize the Sensitivity Requirement and the Closure Principle. Nozick claims, correctly, that no such argument succeeds. But he gets almost all the details wrong. The Sensitivity Requirement and the Closure Principle are compatible; the Sensitivity Requirement is incorrect; and even if true, the Closure Principle is structurally incapable of generating a plausible and valid global skeptical argument. It is therefore a mistake to take the Closure Principle as central in discussions of skepticism. The paper concludes by examining the prospects for a plausible skeptical argument.

Introduction

In the twenty-three years since the publication of *Philosophical Explanations*, the chapters on knowledge and skepticism have maintained a central place in the epistemological discussion. There are three reasons for this. First, Nozick's "tracking" account of knowledge remains one of the most well-developed reliabilist accounts on offer, and so has served as both an exemplar and a stalking horse for a generation of epistemologists. Second, one of the key elements of this account, the "sensitivity" requirement, has been influential in its own right, even apart from the other features of Nozick's view. Finally, the response to epistemological skepticism mooted in these pages remains perhaps the clearest and most vigorous expression of an understanding of skepticism which had been in the air for the previous decade or so and has not lost its allure in the interim.

My primary concern in this paper is the third issue: Nozick's response to skepticism and the understanding of skepticism which underwrites it. As I will argue, his response ultimately fails, and his understanding of the skeptical argument turns out to be incorrect. In particular, like many

other epistemologists over the last 30 or so years, Nozick takes the question of whether knowledge is closed under known implication to be crucial for our understanding of, and response to, external world skepticism. However, as I will try to bring out, this question is a red herring: even if the Closure Principle is true, it does not yield a successful argument for skepticism. This reveals a point of deeper significance. It is often assumed that it is obvious how the skeptic's argument is supposed to work and that the hard questions concern how best to respond. However, this is not so. The pay-off, then, will be a deeper understanding of what it would take to launch a plausible version of skepticism.

Nozick's Response to the Skeptic

Let me begin by reprising a familiar story, though in a somewhat anachronistic way.

Some skeptical arguments are *global* arguments: they aim to show that no one knows anything about the external world because no one knows that certain possibilities—such as that one is dreaming a life-like dream, deceived by an omnipotent evil demon, or a brain in a vat being subjected to appropriate illusory sensory stimulations—do not obtain. Consider, for instance, the following argument, where H stands for the hypothesis that one of these possibilities obtains and O stands for a representative proposition about the external world (e.g., that I have hands, or that there is a computer here):

Argument from Ignorance¹:

1. I do not know that not-H.
2. If I do not know that not-H, then I do not know that O.
3. So, I do not know that O.
4. So, since O is a representative proposition about the external world, I do not know anything about the external world.

If this argument works for my own case, then it equally applies to everyone; it's not something special about *me* that prevents me from having knowledge of the world. So,

5. No one knows anything about the external world.

This argument can be recast so as to be deductively valid.² The premises are tempting, the conclusion patently false. Something has to give, but what?

The best approach, I believe, is to try to understand the appeal of the argument's premises. In particular, what considerations could successfully motivate premises 1 and 2? Though Nozick does not explicitly formulate the Argument from Ignorance, it is instructive to read him as attempting to answer this question. On his view, the skeptic motivates

¹ I owe this label, and the formulation of the argument (through step 3) to DeRose, "Solving the Skeptical Problem," *The Philosophical Review*, 104 (1995), 1–52.

² Just add the premises: A) If I don't know O, then I don't know anything about the world, and B) If I don't know anything about the world, then no one knows anything about the world.

the first premise by noting that if the skeptic's hypothesis *were* correct, one would nonetheless believe that it was not (201). Thus, the skeptic relies on a general principle relating one's possession of knowledge to what one would believe under certain conditions. It can be stated as follows:

Sensitivity: In order for one to know that *p*, it must be the case that: if *p* were not true, then one would not believe that *p*.

And in moving from "I don't know not-H" to "I don't know O", the skeptic "assumes" (204) the following principle:

Closure: For all *p* and all *q*, if one knows that *p* and knows that (*p* entails *q*), then one knows that *q* [assuming one believes it].³

Nozick's text strongly suggests that it is *only* by assuming the Closure Principle that one can successfully motivate premise 2.⁴ This, however, is surely incorrect. Fortunately, we can take him to be making a weaker claim: *one way* to successfully motivate premise 2 is by assuming the truth of the closure principle; if Closure were true, it would generate a requirement capable of doing the work the skeptic demands of it. We can thus see Nozick as exploring the viability of a version of the Argument from Ignorance which supports premise 1 by appealing to Sensitivity and premise 2 by appealing to Closure.

Nozick argues that such an argument fails. We should accept Sensitivity, he claims, but Closure is false. According to his "tracking" conception of knowledge, a person S knows that *p* just if:

1. *p* is true,
2. S believes that *p*,
3. [Sensitivity] if *p* were not true, S would not believe that *p*, and
4. if *p* were true, S would believe that *p*.⁵

If this account is correct, then the Closure Principle is false. There will be cases in which conditions (1)–(4) are met for a given belief *p*, but (3) is not

³ This is only a rough approximation. For my purposes here, further refinements are unnecessary. However, though Nozick is not careful to do so (see, e.g., 205), this principle should be distinguished from a principle of *transmission* which says that if one knows that *p*, knows that (*p* entails *q*), puts two and two together, and infers that *q*, then one can gain knowledge that *q* on the basis of this inference. There may be cases in which the Closure Principle is satisfied, but in which one could not come to know *q* by thus inferring it from *p*.

⁴ He writes, for instance, "The details of different skeptical arguments vary in their structure, but each one will assume some variant of the principle that knowledge is closed under known logical implication" (204).

⁵ Nozick further refines conditions 3 and 4 so that the method of belief formation is held fixed between the actual case and the relevant counterfactual cases. This refinement is motivated by examples of the following sort. Suppose (1) that a grandmother comes to believe, by seeing her grandson, that he is in good health, but (2) if he were sick or dead her relatives would tell her that he was well. Condition 3, as formulated above, is violated, but it seems that the grandmother *does* know her grandson is well when she is looking right at him. A striking feature of the example is that the grandmother's method of belief formation shifts between the actual and the counterfactual case. Nozick accordingly refines condition 3 to hold the method of belief formation fixed:

met for a proposition which S recognizes to be entailed by p . In fact, closure will fail in precisely the cases that concern the skeptic. My belief that I am not a brain in a vat is not sensitive: I would believe that I am not a brain in a vat even if I were a brain in vat. However, my belief that I have hands *does* meet requirement 3. So if I meet the remaining conditions of the tracking account, I can know that I have hands even though (1) I know that my having hands entails that I'm not a brain in a vat, and (2) I do not know that I am not a brain in a vat.

This is not Nozick's only criticism, however. He also holds that Sensitivity and Closure conflict. He writes,

It is clear that *any* account that includes as a necessary condition for knowledge the subjunctive condition 3, $\text{not-}p \rightarrow \text{not-(S believes that } p)$, will have the consequence that knowledge is not closed under known logical implication (207, italics added).

He consequently concludes that the skeptic's position is incoherent, because it depends upon two incompatible principles. As he puts it, "The skeptic cannot be right both times" (209).

This point tells against a common understanding of Nozick's response to the skeptic. It is sometimes worried that Nozick begs the question against the skeptic by assuming that the skeptic's hypotheses do not obtain. The common response on Nozick's behalf is that he is not attempting to refute the skeptic using only claims which the skeptic would accept, but rather showing that we can satisfactorily explain how knowledge is possible, given considerations which *we* accept. This characterization does indeed accord with some of his official pronouncements about his aims in these pages.⁶ However, as should now be apparent, Nozick

"If p weren't true and S were to use [method] M to arrive at a belief whether (or not) p , then S wouldn't believe, via M, that p " (179).

This refinement is irrelevant to our concerns at this point. As explained below, however, the refined requirement is too strong.

Keith DeRose has pointed out that the apparent problem presented by this example can be avoided without explicitly relativizing the requirement to the method of belief formation.

"The standard possible-worlds semantics for counterfactual conditionals can illuminate what's going on here. When one searches for the possible worlds most similar to the actual world in which the grandson is not well, the respects in which the possible worlds are to resemble the actual world is a highly context-sensitive matter. Especially where the context focuses one's attention on the grandmother and her cognitive and recognitional abilities, one *can* place heavy weight upon similarity with respect to the method she is using to arrive at her belief, and then it can seem that in the closest world in which the grandson is not well, she's looking right at him and so does *not* believe he is well" ("Solving the Skeptical Problem," 21).

DeRose accordingly suggests that we could leave the requirement as originally formulated, keeping in mind that, ordinarily, when we are evaluating the truth of the relevant subjunctive conditionals, heavy emphasis will appropriately be placed upon the person's method of belief formation. The flexibility thus introduced into our understanding of these conditions is, to my mind at least, a virtue. In what follows, I will assume this interpretation unless I explicitly state otherwise.

⁶ See, e.g., 197-8.

also tries to do more. He charges that the envisaged route to the skeptical conclusion saddles one with an inconsistent position.

Sensitivity and the Closure Principle: Why Nozick's Response to the Skeptic Fails

It should be granted that *if* Nozick's account of knowledge is correct, then (1) there are cases in which the Closure Principle is violated and (2) some of the skeptic's hypotheses will be among those cases. However, Nozick's response to the skeptic fails on every other point.

First, the Sensitivity Requirement is not incompatible with the Closure Principle.

Here is Nozick's argument to the contrary:

The further conditions [beyond true belief] will make knowledge open under known logical implication ... when at least one of the further conditions itself is open. Knowledge stays closed (only) if all of the additional conditions are closed (208).

This is not correct. The Sensitivity Requirement just says that in order for one to know that p , it must be the case that if p were not true, one would not believe that p . This is simply a *necessary* requirement for knowledge. One could hold a view which endorsed this necessary requirement and supplemented it with further requirements which (jointly) had the consequence that one cannot know anything at all. Such a view could also endorse the Closure Principle. For that principle will be trivially satisfied if one *never* knows any of the propositions which might be substituted for p and q . Hence, Sensitivity and the Closure Principle are logically consistent. Nozick's skeptic is not vulnerable to the charge of inconsistency.⁷

Is it pragmatically incoherent to deploy both Sensitivity and Closure in the way that Nozick's skeptic deploys them? Again, no. To see this, consider an example. Imagine a skeptic who endorses Sensitivity and consequently maintains that in order for one to have knowledge that p , one must not have a false belief that p in the nearest $\text{not-}p$ worlds. This skeptic notes that we do not meet this requirement when $p = I \text{ am not a brain in a vat}$. He concludes that we do not know the truth of this proposition. However, this skeptic also endorses the closure principle. So, by modus tollens, this skeptic concludes that one does not know such mundane things as that one has hands.⁸ This skeptic is not in a pragmatically incoherent position. All that has happened is that his commitment to the closure principle has the consequence that having a true, sensitive belief does not *suffice* for knowledge. But he never said that it did. His view is that satisfying Sensitivity is *necessary* for possessing knowledge—and that for that reason we cannot know that we are not brains in vats—but on his view it

⁷ Cf., Williamson, *Knowledge and Its Limits* (Oxford: University Press, 2000), chapter 7. Nozick's claim here is echoed by Klein, "Skepticism," in the *Oxford Handbook of Epistemology*, Paul Moser, ed. (Oxford: University Press, 2002), 347.

⁸ This skeptic is not merely an *ad hoc* creation. It is, more or less, the "high-standards" skeptic who figures centrally in DeRose's "Solving the Skeptical Problem."

turns out that in many cases a whole lot more than sensitivity is required for knowledge as well.

Consider now Nozick's charge that the Closure Principle is false. This charge may be correct. (I take no stand on that issue here.) However, Nozick does not offer a convincing argument for it. His rejection of the Closure Principle depends upon his analysis of knowledge. But his analysis of knowledge is vulnerable to persuasive counterexamples because of its commitment to the Sensitivity Requirement. Here is a perfectly straightforward one. Right now I know that George W. Bush is President of the United States. However, my belief that George W. Bush is President does not meet condition 3 (Sensitivity). If George W. Bush were no longer President, it would be because he had just had a heart attack, or had just been assassinated, or had just died in a plane or helicopter wreck, or some such thing. In all of those cases, I would still believe that he was President, since there is a considerable time-lag between events relevant to Bush's status as President and my awareness of them. But this does not prevent me now from knowing that he is the President.⁹

⁹ This counterexample was inspired by John McDowell's illuminating discussion of a similar example in "Knowledge by Hearsay," in McDowell, *Meaning, Knowledge, and Reality* (Cambridge, MA: Harvard University Press, 1998), 414–443. McDowell does not discuss the bearing of such examples on the Sensitivity Requirement.

A number of putative counterexamples to the Sensitivity Requirement have been discussed in the literature. Most founder on delicate issues relating to the evaluation of subjunctive conditionals and relativization to methods of belief formation. For example, consider the following example from Peter Klein.

You put a glass of ice-cold lemonade on a picnic table in your backyard. You go inside and get a telephone call ... and talk for half an hour. When you hang up you remember that you had left the ice-cold lemonade outside exposed to the hot sun and come to believe that it isn't ice-cold anymore. It would seem that you could know that, even if in some near world a friend of yours who just happened to be walking by noticed the glass and, happening to have a cooler full of ice with him, put the glass of lemonade in the cooler to keep it ice-cold for you. Thus, if the lemonade were still ice-cold, you would believe that it wasn't" (Klein, "Skepticism," 346, derived from an example from Vogel, "Tracking, Closure, and Inductive Knowledge", in Luper-Foy, ed., *The Possibility of Knowledge: Nozick and his critics* (Oxford: Rowman and Littlefield, 1987)).

This example runs into trouble because of the vague specification "in some near world...". If the described world is the *nearest* world in which the lemonade is still ice-cold, then the belief is not sensitive. But by the same token, if one feels the tug of the sensitivity requirement, then it is not obvious that in such conditions one *does* know that the lemonade isn't ice cold. On the other hand, if we suppose the ordinary sort of conditions to prevail, then it is implausible that the nearest world in which the lemonade is still ice-cold is one in which, unbeknownst to you, your friend happened to walk by with a cooler full of ice and, recognizing your glass, kindly put it in the cooler for you. Rather, the nearest world in which your lemonade is still ice-cold is (arguably) the world in which you didn't forget to bring your lemonade with you when you went inside to get the phone. In this latter world, of course, you don't believe that your lemonade is not ice-cold. So your belief is sensitive, and the example—so conceived—is not a counterexample after all. (As I'll explain shortly, this response is in accord with Nozick's theory of the evaluation of subjunctive conditionals.)

However, examples like this one *do* tell against Nozick's official view, since his version of Sensitivity requires that the method of belief formation be held constant. On his final statement of the Sensitivity requirement, it reads: "If *p* weren't true and S

Let's consider this example more carefully. Nothing here trades on considerations about "methods" of belief formation; the counterexample

were to use [method] M to arrive at a belief whether or not *p*, then S wouldn't believe, via M, that *p*" (179). This means that to determine whether S knows that *p*, we must look at the nearest world in which *p* is false and S forms a belief about the matter utilizing the method used in the actual world. In the above example, the method used is something like memory (of leaving the lemonade outside) plus inductive inference (from one's knowledge of what happens to cold things which are left in the sun). In the nearest world in which one forms a belief about the state of one's lemonade using that method and the lemonade is still cold, one would get it wrong. However, it seems that one does know that the lemonade is no longer ice cold. This suggests that Nozick's official, revised version of the requirement is too strong.

Jonathan Vogel suggests amending Nozick's view by removing the reference to methods from the antecedent of the subjunctive conditional, as follows:

If *p* were false, it would not be the case that S would believe that *p* via method M (where M is the method S actually uses [i.e., the method S uses in the actual world] ("Tracking, Closure, and Inductive Knowledge", in Luper-Foy, ed., *The Problem of Knowledge*, at 198).

This requirement is very weak, since it is satisfied—regardless of whether S would believe *p*—so long as S would have different evidence or use different methods if *p* were false (op. cit., 199–200). To my mind, it evades counterexamples such as the example of the ice-cold lemonade, since the closest world in which the lemonade is still cold is the world in which you brought it in out of the sun—and in that world, your belief as to whether the lemonade is ice-cold would have arisen from different evidence. Vogel, however, offers a similar example as a counterexample even to this weak version of the Sensitivity Requirement (206). In doing so, he follows David Lewis' semantics for subjunctive conditionals, which holds that the most similar world is one whose history is identical to that of the actual world up until the point when the ice cubes began to melt. One's evidence in that world would be the same as one's evidence in the actual world, and one would believe the relevant *p* even though it is false. However, as Vogel acknowledges, Nozick himself rejects this aspect of Lewis' view, commenting that "the relevant not-*p* world is not a world identical to the actual one until now, and then diverging so as to produce not-*p*..." (*Philosophical Explanations*, 223n). In particular, Nozick permits so-called "backtracking compound" conditionals, in which a subjunctive conditional whose antecedent refers to a time earlier or identical to the time referred to in the consequent is evaluated in light of a subjunctive conditional whose consequent states how things *would have to have been* in order for the antecedent of the first conditional to come about. For instance, consider the conditional:

If the lemonade were still ice-cold, I would not believe as a result of memory and induction that it was not ice-cold.

On Nozick's view, unlike Lewis's, this conditional comes out true. For if the lemonade were still ice-cold, it would be because I had brought it in with me. And if I had brought it in with me, I would not believe on the basis of memory and induction that it was no longer ice cold. Consequently, Nozick would not regard the example as a counterexample to Vogel's revised version of the Sensitivity requirement. (The counterexample which I offer in the main text is superior in this regard.)

Ernest Sosa offers a counterexample similar in certain respects to mine in "Tracking, Knowledge, and Competence", in *The Oxford Handbook of Epistemology*, 264–286. As he notes, however, his example, too, raises issues relating to the use of backtracking compound conditionals. A significant advantage of the counterexample presented in the main text above is that it can be cast in a way that skirts these technical issues. Backtracking conditionals are not required to state it, but even if they are permitted, Nozick could not plausibly evade the counterexample by appealing to them. The example is consequently neutral in the disagreement between Lewis and Nozick over the semantics of subjunctive conditionals.

succeeds even if we incorporate Nozick's suggestion that the method of belief formation should be held fixed between the actual case and the not- p cases under consideration. Whatever the "method" is by which I now believe that Bush is President, in a close world in which Bush were no longer President, that "method" would yield a false belief. It would only be when I resorted to further sources of information (e.g., listening to the news, or chatting with colleagues) that I would come to have a true belief about the matter. And even if we characterize the method so that it includes such things as periodically checking the news, the belief still will not be sensitive. Even if I were to spend all my waking hours listening to news radio solely in order to gain correct beliefs about who is President, I would still have false beliefs about the matter for some non-negligible time period if Bush were to cease to be President. Yet when Bush *is* President and nothing untoward has gone on, I know it.

It might be suggested that what matters here is that my belief is responsive, over time, to the facts. I engage in a diachronic belief-forming policy or procedure of regularly checking the news and participating in ordinary conversations, so if Bush were to cease being President, I would be alerted to that fact and so would cease to believe that he was President within a reasonable period of time. It is tempting to think that if this were not so, then I would not now know that Bush is President. (Consider someone who followed the policy of checking the news only once every election cycle.) So we might try revising the Sensitivity Requirement to accommodate temporal slack, thus:

In order for one to know p , it must be the case that: If p were false, then one would not believe, or would cease to believe, that p *within some appropriate time period*.¹⁰

However, this requirement, too, is incorrect. Suppose, for instance, that I have decided to enter a monastery where I will be completely isolated for the remainder of my life from anyone or anything that could inform me of worldly affairs. I enter the monastery knowing that George W. Bush is President. (I read the papers and listened to the news that morning and participated in ordinary conversations on the way to the monastery.) But if Bush weren't then President, I would never learn of that fact. My beliefs would not be revised within an appropriate time period. Of course, given my background knowledge, I would at some point cease believing that he is President. But this is not the sort of time-delayed sensitivity or responsiveness to the facts that the proposal is meant to capture. Thus time-delayed sensitivity or responsiveness to the facts is not required for knowledge. It is worth noting, incidentally, that this is not an isolated case. Much of our everyday knowledge does not involve a policy or method of belief formation involving periodic "checking up" so as to ensure periodic adjustments to the facts. Consider, for instance, your knowledge of such things as that the remains of last night's dinner are in the trashcan.

¹⁰ What is to count as an "appropriate" period of time may, for our purposes here, be left as a contextually-determined matter.

Keith DeRose has defended something very much like the Sensitivity Requirement, not as a necessary condition for knowledge, but as a general characterization of one consideration which will prompt us to judge that someone does not have knowledge. As he puts it, "We tend to judge that S doesn't know that P when we think that S's belief that P is insensitive".¹¹ DeRose considers certain apparent counterexamples to this claim. For instance, consider my belief that I don't falsely believe that I have hands. If I *did* falsely believe that I have hands, I would still believe that I have hands (by hypothesis) and so would still believe that I do not falsely believe that I have hands. Hence, this belief is not sensitive. But don't we know that we do not falsely believe that we have hands—or at least, aren't we inclined to so judge? DeRose takes such examples to indicate a certain limitation on the generalization about our patterns of knowledge ascription: "We *don't* so judge ourselves ignorant of P where not- P implies something we take ourselves to know to be false, without providing an explanation of how we came to falsely believe this thing we think we know."¹² However, this proposal does not help with our example. Given that we haven't yet heard the news, Bush's no longer being President provides the needed explanation. In response to certain problems with DeRose's proposal, Timothy Williamson suggests that we refine it as follows: "when we judge [Sensitivity] false, we do so because S sensitively believes a proposition q which entails p , and $\sim p$ does not explain how S could falsely believe q ."¹³ However, this suggestion does not account for our example either. For what could be the relevant proposition q which entails that George W. Bush is President? The only plausible contender is that I remember that George W. Bush is President. But is this belief—that I remember that George W. Bush is President—itself sensitive? It is hard to say. Much here depends upon how we characterize the method by which this belief is formed. If we characterize the method internally, as involving the fact that *I seem to remember that George W. Bush is President*, then it seems that this belief is not sensitive: in the nearest world in which the belief is false (e.g., because Bush is no longer President) but I

¹¹ "Solving the Skeptical Problem", 18. DeRose takes this point to support the "Rule of Sensitivity" which forms the basis for his contextualist response to the Argument from Ignorance:

Rule of Sensitivity: When it's asserted that S knows (or doesn't know) that P, then, if necessary, enlarge the sphere of epistemically relevant worlds so that it at least includes the closest worlds in which P is false ("Solving the Skeptical Problem," 37).

That is, the assertion that S knows (or doesn't know) that P raises, or tends to raise, the conversationally-relevant standards for knowledge so as to require that one's belief that P be sensitive *and* that one's other beliefs track the truth in the nearest not- P world as well as in all closer counterfactual worlds. DeRose's Rule of Sensitivity is vulnerable to the same example which felled Nozick's view; if I claim to know that Bush is President, no one will be inclined to dispute my claim on the grounds that this belief is insensitive. In fact, no one in an ordinary conversational setting (outside of a discussion of skepticism) would be inclined to dispute it at all.

¹² "Solving the Skeptical Problem," 23.

¹³ Knowledge and Its Limits, 159. Williamson does not endorse the proposal.

form a belief using this method, I will believe, incorrectly, that I remember that Bush is President. Moreover, even if this belief is regarded as sensitive, the explanatory requirement is met. For George W. Bush's no longer being President does, nearly enough, explain how I could falsely believe that I remember that George W. Bush is President—if, for instance, I have not yet heard the news.

These considerations strongly suggest that the Sensitivity Requirement is wrong in spirit, both as a necessary condition for knowledge and as a generalization about our practice of knowledge attribution. There are central cases of knowledge in which we do not meet it or anything like it, and we are not inclined to impose it even when we explicitly reflect on the question of whether we possess knowledge in these cases.¹⁴

It might seem that rejecting the Sensitivity Requirement amounts to a repudiation of skepticism, since according to Nozick's analysis the skeptic, too, appeals to the Sensitivity Requirement in order to argue that we do not know that the skeptical possibilities don't obtain. However, this analysis is incorrect as well. A traditional skeptic appeals to considerations about your *reasons for believing* that the skeptical possibilities don't obtain, arguing that your evidence does not give you an adequate reason to believe as you do. This sort of argument requires, of course, that the skeptic is working within an evidentialist conception of knowledge or justified belief, and that the skeptic can somehow motivate the thought that the relevant evidence in these cases is confined to how things *experientially appear* to us. Both requirements raise substantial difficulties. But the important point here is that once this framework is in place, the skeptic can argue, for instance, that *we don't have adequate evidence that we are not just asleep and dreaming* and consequently *don't have a good reason to believe that this possibility does not obtain*. This charge is distinct from the charge that our belief is insensitive. For it might be the case that one *has good reason to believe* that *p*, and yet if *p* were false, one would still believe that *p*. (To see this, just let *p* = George W. Bush is President.)¹⁵

In sum, Nozick claims: 1) Closure and Sensitivity are incompatible; 2) The tracking account of knowledge entails the denial of Closure; 3) The tracking account of knowledge is correct; 4) Sensitivity is correct; 5) The

¹⁴ What it should be replaced with is a difficult matter (and a topic for another paper). For recent discussion, see for example Sosa and Williamson.

¹⁵ Mark Kaplan has pointed out to me that while Stroud does not appeal to Sensitivity in his account of why we don't know that we aren't dreaming, he *does* say that perceptually knowing that *p* is incompatible with dreaming that *p* because if one came to believe that *p* as a result of a dream and *p* were also true, the truth of one's belief would be a mere "coincidence" (*The Significance of Philosophical Scepticism* (Oxford: Oxford University Press, 1984), chapter 1). This perhaps suggests a commitment to the Sensitivity Requirement. However, there are surely other ways of explaining why dreaming is not a way of gaining knowledge. For instance, one straightforward possibility would be to suggest that given the way the world works, dreams are not objectively good evidence of how things are. Or one might suggest that beliefs resulting from dreams fail to satisfy some reliability requirement other than Sensitivity.

skeptic must presuppose Sensitivity; 6) The Closure Principle can be exploited to yield a requirement playing the role of premise (2) in the Argument from Ignorance. I have argued that claims (1), (3), (4), and (5) are incorrect. (2) should be granted. What of (6)? What role can the Closure Principle play in the Argument from Ignorance?

The Closure Principle Does Not Generate a Compelling Skeptical Argument

Nozick is not alone in holding that exploiting the Closure Principle is *one* way to generate a plausible skeptical argument. It is widely assumed that the Closure Principle can underwrite premise 2 of the Argument from Ignorance. However, as I will argue in this section, this assumption is incorrect. Even if the Closure Principle is true, it is structurally incapable of yielding a plausible argument for external world skepticism.

It will be helpful, first, to be clear about what is at issue here. The global skeptic's claim is that *no one* can know *anything* about the world outside of one's own mind. The skeptic supports this claim in part by maintaining that one can't know anything about the world unless one knows, e.g., that one is not dreaming. The question, then, is this: can the closure principle generate a version of this requirement which is capable of providing a compelling argument for the global skeptical conclusion?

It cannot. To see why, let's begin with a fundamental point. The possibilities that one is asleep and dreaming a life-like dream, that one is a brain in a vat whose experiences are the result of stimulations from a supercomputer, and that one's experiences are being controlled by an all-powerful demon are all *perfectly compatible* with the truth of the vast majority of one's beliefs about the world, even the truth of one's beliefs about the location and features of objects in one's immediate vicinity. For instance, the possibility that my experiences are being caused in one of these ways is perfectly compatible with the truth of my belief that there is a table here. Consequently, the entailment upon which the Closure Principle depends simply does not hold in the vast majority of cases. If the skeptic's requirement is underwritten by the closure principle, then skepticism is true *at most* for a small portion of our beliefs about the world.¹⁶ (It should be noted that the skeptic's hypotheses *are* incompatible with one's having perceptual knowledge that *p*. For instance, one can't know that *p* if one merely came to believe it as the result of a dream; knowing that *p* entails that one did not merely dream that *p*. However, such entailments cannot be exploited by the Closure Principle, since it concerns the implications of what is known, not of one's knowing it. We'll return to the question of whether the latter implications enable some-

¹⁶ This point is also stressed by Stroud, *The Significance of Philosophical Scepticism*, chapter 1. Of course, the Closure Principle might figure in a skeptical argument which depended on the thought that for any belief about the world, there is some incompatible possibility which is not known not to obtain. But such an argument is not the global version of the Argument from Ignorance which is under investigation here.

thing very much like the Closure Principle to play a role in skeptical argumentation.)

It might be thought that the above problem could be removed by characterizing the skeptic's possibility so that it is incompatible with the truth of one's beliefs about the world quite generally. For instance, perhaps the relevant possibility is that one is *falsely dreaming*, or that one is being *deceived* by a mad scientist or Evil Demon. However, there are two problems with this move. First, it doesn't adequately account for the intuitions that fuel the skeptic's requirement in the first place. Imagine, for instance, that you are an Evil Demon's plaything: for a good amount of the time he deceives you, but sometimes he doesn't, and you can't tell the difference. Being an Evil Demon's plaything is just as epistemically disastrous as always being the victim of an Evil Demon's deceptions, and a skeptical argument based on the possibility that you are a plaything would have just as much force as an argument based on the possibility that you are a victim of constant deception. Likewise, consider that if you come to be convinced of the truth of p merely as the result of *dreaming that p* , then you don't know that p —even if p is true. A skeptical argument based upon this latter possibility has a great deal of intuitive force. Thus, no analysis which, like the closure principle analysis, requires the skeptic's possibility to be incompatible with the truth of our beliefs will yield a satisfactory explanation of the appeal of the skeptic's requirement.

The second—and more important—problem with the proposal is that it renders the skeptic's possibilities idle: they no longer play an important role either in generating the skeptic's requirement or in the argument that one does not know that p . To see this, consider how an instance of the skeptic's requirement is generated for a belief whose truth is compatible with the possibilities that one is dreaming and that one's experiences are being caused by a scientist or demon. On the suggested account, the relevant skeptical hypothesis can be understood as combining two distinct components: (1) the hypothesis that one is being led to believe that p by a life-like dream, an all-powerful demon, or a scientist electrically stimulating one's brain, and (2) the hypothesis that p is false. A proposed explanation is then added for the hypothesized state of affairs: the demon or scientist is aiming to deceive. The incompatibility between this hypothesized state of affairs and most propositions about the world arises only from the second hypothesis, "not- p ." So though a relevant instance of the closure principle is generated, it's only because p implies not-not- p . What enables the skeptic's possibility (e.g., That one is dreaming that p) to appear in the relevant instance of the closure principle is just that p also implies not-not- $(p \vee q)$. To see how this works, let q be *I am not dreaming that p* . Then we have an instance of the closure principle (removing the double negations):

If I know that p and know that p implies $(p \vee I \text{ am not dreaming that } p)$, then I know that $(p \vee I \text{ am not dreaming that } p)$.

This is equivalent to:

If I know that p and know that p implies not(not- p and *I am dreaming that p*), then I know that not(not- p and *I am dreaming that p*).

So now, given our knowledge of the relevant implication, we have the skeptical requirement:

If I know that p , then I know that it is not the case that: I am dreaming that p and p is false.

That is:

If I know that p , then I know that it is not the case that: I am falsely dreaming that p .

However, in this route to the skeptical requirement, the possibility that I am dreaming that p just came along for the ride. Any other proposition could have been substituted in its place. Consequently, the *particular content* of the skeptical possibility isn't playing any role at all in generating the requirement.

Consider, now, the skeptic's claim that this requirement is not met. To meet the requirement, one must know that it is not the case that: one is dreaming that p and p is false. One could know this, and hence meet this requirement, by knowing *either one* of two things: either that one is not dreaming, or that p is true.¹⁷ So the skeptic's argument that one does not meet this requirement must have two parts. First, there must be a sub-argument that one does not know that one is not dreaming that p . Second, another sub-argument is needed showing that one does not know that p . But that is just what the skeptical argument was supposed to show in the first place. So if the skeptical hypotheses are characterized in the proposed way, then even a successful argument that one does not know that one's experiences or beliefs are not being generated in one of the deviant ways will fail to establish the skeptical conclusion. For instance, even if the skeptic demonstrates that one does not know that one is not dreaming (where this possibility is compatible with the truth of the belief in question), one could still meet the requirement generated by the closure principle *by knowing that it is not the case that p is false*. The skeptic still needs some other argument that one does not know that p .

The discussion so far has treated the skeptical hypothesis as a simple conjunction of 'I'm dreaming' and 'not- p '. It might plausibly be suggested that the skeptic's hypothesis is rather the more complex explanatory hypothesis, 'I falsely believe p because I am dreaming.' However, this suggestion does not evade the objection. The suggested hypothesis amounts to this: I have a false belief whose falsity is to be explained by the fact that I am dreaming. However, while the falsity of one of my beliefs could be explained by the fact that I am dreaming, this explanatory relation does not arise from an entailment, since dreaming that p is compatible with the truth of p . For this reason, the explanatory appeal to dreaming in the

¹⁷ It might be suggested that it would be sufficient for one to know that $(p$ or one is not dreaming that $p)$. This is correct, truth-functionally speaking. But I do not see how one could know the truth of this disjunction without either knowing that p or knowing that one is not dreaming that p .

proposed skeptical hypothesis doesn't generate an instance of the closure principle; it is the falsity of p that does it. Thus, as before, an instance of the closure principle is generated merely because of the hypothesized falsehood of p ; the explanatory appeal to dreaming just comes along for the ride. Likewise, one could know that the complex explanatory hypothesis does not obtain *either* by knowing that p is true *or* by knowing that one's belief was not caused by a dream. So, again, even if the skeptic can show that one cannot know that one's belief was not caused by a dream, the skeptic will still need some other argument that one does not know that p .¹⁸

It is tempting at this point to object like this: "If the skeptic has successfully argued that I can't know that I am not dreaming, then hasn't he undermined my knowledge that p ? How could I know that p if I don't know that I am not dreaming?" In offering this reply, however, one must be appealing to some other version of the skeptic's requirement, *not to the one that is generated by the closure principle*, since the requirement generated by the closure principle does not appropriately tie together one's knowledge that p and one's knowledge that one is not dreaming. But if some other version of the skeptic's requirement is in play, then the closure principle is otiose. (And our inclination to offer this response at this point indicates that the closure principle is not what moves us in the first place to accept the skeptic's attempt to link our possession of knowledge of the world to our ability to know that we are not dreaming.)

An advocate of the closure principle account might explore two moves at this point. The first is to grant the above argument, but to suggest that the closure principle operates in tandem with the so-called KK principle, that if one knows that p , then one knows that one knows it. Even if the skeptic's possibilities are compatible with the truth of the belief in question, they are not compatible with one's possessing *knowledge* (as I earlier noted). Consequently, if the KK-principle is true, one can substitute " Kp " for " p " in the closure principle, thereby generating a principle with the requisite generality. The trouble with this move, however, is that the KK-principle is false.¹⁹ A skeptical argument which requires it would not be worth much attention.

A second way of saving the closure principle analysis would be to enrich the Closure Principle so as to exploit the incompatibility between the skeptic's possibilities and knowing that p , as follows:

If you know that *knowing* that p implies the truth of q , and if you know that p , then you also know that q .²⁰

This enriched version of the closure principle would yield a skeptical argument which applies quite generally to one's beliefs about the world, but there is a simple reason why it can't do the work it is being asked to do. The skeptical argument is supposed to yield the conclusion that *no*

one (at least no human being) knows anything about the world. But any plausible closure principle will contain the clause "if one *knows* that implies," and a skeptical argument based upon it consequently will not apply to someone who does not know about or does not believe the relevant implication. But this result is ludicrous. Ignorance is no way to evade the truth of skepticism, if it is true. The skeptical requirement must be accounted for in a way which explains why it applies to everyone if it applies to anyone.

It might be suggested that in order even to count as a knower, someone would have to recognize the relevant implication. But this is surely incorrect. For instance, there have been, and probably still are, many people who think that dreaming that p is a way of gaining knowledge that p . Their possession of this incorrect belief would not, by itself, preclude them from counting as knowers. Moreover, for *every* possibility capable of doing the requisite work in the enriched closure principle (e.g., that one is a brain in a vat, or that one is an evil demon's plaything), someone could fail to know about the incompatibility between coming to know something about the world and that possibility's obtaining, and this failure would not, by itself, preclude him or her from being a knower. Some of my students, for instance, learn something new, something which they did not know before, when I teach them about the incompatibility between the skeptic's possibilities and their possession of perceptual knowledge about the world.²¹

It should go without saying that eliminating the requirement that the implication be *known* results in a blatantly false requirement. The resulting principle would have the result that if you know anything, then you know all logical truths. The challenge, then, if we are to successfully motivate the second premise of the Argument from Ignorance, is to explain why *no one* can possess knowledge of the world—regardless of what they happen to know or believe—unless they know that the skeptic's possibilities don't obtain, and to do so in a way that does not commit us to absurdly strong requirements.

Given these considerations, I conclude that even if the Closure Principle is true, it will not generate a plausible and valid version of the Argument from Ignorance. In a sense, then, I agree with Nozick: a global skeptical argument will not succeed if it depends upon Sensitivity and Closure. But unlike Nozick, I do not endorse this conclusion on the basis of a contested theory of knowledge. Rather, I endorse it because (1) the Sensitivity Requirement is false and (2) the Closure Principle cannot do the work it is being asked to do.

Postlude

There is one further problem with the claim that the second premise of the Argument from Ignorance can be motivated by the Closure Principle. If the second premise is motivated merely by the Closure Principle, then

²¹ I am grateful to Mark Kaplan for raising the issues addressed in this paragraph.

¹⁸ I am grateful to an anonymous referee for raising this objection.

¹⁹ For a recent provocative discussion, see Williamson, *Knowledge and Its Limits*, chapters 4 and 5.

²⁰ For the suggestion that skepticism involves such a requirement, see Stroud, *The Significance of Philosophical Scepticism*, chapter 1.

(given the falsity of Sensitivity) it is hard to see why we should think that one does not know that the skeptic's hypotheses don't obtain. As G. E. Moore in effect suggested, why can't you meet this requirement just by knowing some appropriate p , recognizing that p implies that you aren't dreaming, and believing on this basis that you aren't dreaming?²² In order to preclude this response the skeptic needs, minimally, the requirement that in order to know any given proposition p about the world, you must know *on the basis of something other than p* that you are not dreaming that p . In fact, an adequate motivation of the skeptic's requirement would have to entail more: that you cannot appeal to any other claims about the world whatsoever in order to establish that you are not dreaming. For if you can exclude the skeptic's possibilities by appealing to claims about the world, then the skeptical argument will never get off the ground. But no closure principle alone will generate such a requirement. This means—at a minimum—that if the closure principle were to play any role in the Argument from Ignorance, then some other requirement would have to be in play as well.

We can now see more clearly what it would take to motivate the second premise of the Argument from Ignorance in a way that is adequate to the skeptic's purposes. First, an adequate motivation would generate a requirement applying to one's knowledge of the world quite generally: it would have the consequence that knowing any given proposition about the world requires knowing that the skeptical hypotheses don't obtain, even if those hypotheses don't entail the falsity of one's beliefs about the world. Second, it would generate a requirement applying to everyone, regardless of what they happen to know or believe, but without being absurdly strong. Finally, it would prevent one from appealing to any other beliefs about the world in order to establish that the skeptical hypotheses don't obtain. This last requirement is the kicker. For without it, the skeptic will not be able to overcome the fact that it is generally quite legitimate to appeal to other propositions about the world in order to justify or establish the truth of any given proposition about the world.²³ Our response to the Argument from Ignorance therefore hinges on this question: can the skeptic's second premise be motivated in a way that meets all three of these requirements? If not, then we will be free to respond, as common sense would have us do, by granting the second premise but maintaining that we do know that we are not asleep and dreaming, brains in vats, or victims of an evil demon. That, I think, would be a fully satisfying response.²⁴

²² See "Certainty," in G. E. Moore, *Selected Writings* (London: Unwin, 1953), 190.

²³ This point has been emphasized, in different ways, by Barry Stroud and Michael Williams.

²⁴ I would like to thank Mark Kaplan, Ram Neta, and Jim Pryor for helpful discussions relating to this paper.