

AN EMPIRICAL RESPONSE TO EXTERNAL WORLD SKEPTICISM

Adam Leite (Indiana University, Bloomington)

Pacific APA, April 2009

Draft of March 1, 2009

My topic today is external world skepticism. More precisely, it is arguments that utilize certain hypotheses – that one is dreaming, deceived by an evil demon, or a brain in a vat – to argue at one fell swoop that one does not know, is not justified in believing, or ought not to believe most if not all of what one currently believes about the world around one. More precisely yet, my topic is the relation between such arguments and ordinary life.

A good part of the interest and mystique of discussions of skepticism over the last forty years stems from the contention that the seeds of the skeptical conclusion lie in ordinary life: in the commitments, principles, and requirements that inform our ordinary practices of epistemic evaluation, of justifying our beliefs, and of deciding what to believe [ref: Stroud, Clarke, Cavell]. I doubt, however, that this is really so. To explain why, I will approach skeptical arguments from a certain vantage point. I will try to stand, as far as possible, with both feet squarely in ordinary life. I will start out with all of our ordinary commitments about what is the case, about what we know and what we have reason to believe, about when someone knows, is justified, or has good reason to believe something, and about how one should proceed in deciding what to believe. My question, then, will be this. Can I somehow be moved in a reasonable way from that vantage point to accept the conclusion that I know or am justified in believing far less about the world around me than I thought, or that epistemically speaking, I really ought not believe much of what I have believed about the world around me? In order for such movement to take place, I will have to find reason from within my ordinary standpoint to discount or reject much of what I ordinarily accept. And of course, that reason will have to come from some of what I accept from within that ordinary position. So if such movement were to take place, it would show that our ordinary standpoint is interestingly incoherent. At the end of the day, I think, no such movement takes place: there is no reasonable path from our ordinary position to the skeptical conclusion.¹ What I will do today is take on one important piece of this project. I'll press just one line of thought with the hope of showing that getting a reasonable argument going for the skeptical conclusion is a lot more difficult than is often assumed. Some of the detail of what I will be saying has already been noted in the literature. To my knowledge, however, no one has put all the pieces together. I think the basic point I want to urge is one whose time has come.²

The skeptical argument that I wish to explore can be quickly stated in a form that has come to be known as “The Argument from Ignorance”.³

1. I can't know (reasonably or justifiably believe) anything about the world around me (e.g., that there is a piece of paper here in front of me) unless I know (reasonably or justifiably believe) that I am not being deceived by an evil demon (not dreaming, not a BIV being fed misleading inputs).

2. But I don't know (reasonably or justifiably believe) that I am not being deceived by an evil demon (not dreaming, not a BIV being fed misleading inputs).
3. So, I don't know (reasonably or justifiably believe) anything about the world around me (including that there is a piece of paper here in front of me).

The argument can also be recast as a piece of first-person critical deliberation about what to believe.⁴

Test claim for scrutiny: Here's a piece of paper before me.

Objection: But here's a possible scenario: I am the victim of an all-powerful evil demon who deceives me about what is going on around me.

Principle: If I can't provide an adequate response to this objection, in the form of some consideration that tells against this scenario, then I should not hold this belief.

Recognition: I can't provide an adequate response to this objection: There is nothing that I can appeal to against this possibility.

Conclusion: So, I should not believe that there is a piece of paper before me.

Generalization: This test case is just a stand-in for all the things I currently believe about the world around me on the basis of the senses. The same reasoning applies to each of them. So I should not believe any of those things.

These are relatively crude presentations of the arguments, and there are lots of issues about which one might fuss. But I'm presenting matters this way so that I can easily highlight what I think is the fundamental issue here: the second premise of the Argument from Ignorance and the moment (which I'm calling "Recognition") at which one supposedly finds that there is nothing that one can appeal to in response to the suggestion that one is being deceived by an evil demon. How is one supposed to find oneself in agreement with these claims?

Consider the position we were starting from when we encountered these arguments. We started out with lots of convictions about the world and about our epistemic position in relation to it. The reasoning was supposed to *move us* from our ordinary position to the skeptical conclusion. But I cannot see how that has happened. Consider the Argument from Ignorance. Here are some of the things that I believed when the argument started. There are no evil demons. There are no brains in vats – not yet, at any rate. And dreams have certain distinctive characteristics and features which my current experience lacks, so my current experience is nothing like a dream.⁵ Moreover, I started out taking myself to *know* these things and to have *lots of good reasons for believing them*. I could supply some of those reasons if asked to do so. So starting from my ordinary position, it looks as if I should simply reject the second premise of the Argument from Ignorance: I know that I'm not being deceived by an evil demon, because there aren't any evil demons. Likewise, it seems that starting from my ordinary position, I ought to reject the step in the critical deliberation that I labeled

“Recognition”. Of course there’s something I can appeal to against this hypothesis. It has nothing whatsoever in its favor. There are no evil demons. Those look to me like considerations that warrant dismissing the suggested hypothesis as ridiculous.

Now I’m a serious epistemologist, and I’ve been thinking about skepticism for a very long time. I know that I’m not supposed to respond to the skeptical arguments in these ways. But I want to know why not. Why can’t I make use of considerations about the world in responding to these arguments? All the crucial work must actually be done here: in explaining what’s wrong with these sorts of responses. Do any principles or commitments of our ordinary epistemic practices preclude them? Taking the evil demon hypothesis as my stalking horse, I will argue that they do not. I will take for granted that our ordinary epistemic commitments are *fallibilist*, so that a skeptical argument based on a demand for absolute certainty will not show how the seeds of skepticism could lie in ordinary life [ref. Stroud, Leite 2004]. I will focus on issues having to do with the *adequacy of our evidence* and the related charge that the reply I envision is somehow *objectionably question-begging*. One point that I will try to bring out is that the deeper we press into our ordinary practices in these regards, the more resistant they look to the kinds of interpretations that would license the skeptical conclusion.⁶

Begging the question, first pass

Recall that I proposed to reject the second premise of the Argument from Ignorance by noting that there are no evil demons. Let’s begin with a simplistic version of the charge that this reply is objectionably question-begging. “The skeptic claims that you don’t know or justifiably believe anything about the world. In making use of claims about the world in responding to the skeptic, you have to represent or take yourself to know or at least justifiably believe them. So in responding to the skeptical argument in the way you suggest, you are unacceptably begging the question against the skeptic: you are assuming or presupposing the truth of the very thing the skeptic denies.”

I’m willing to grant that my proposed response to the skeptical argument in some sense assumes or presupposes the truth of the very thing the skeptic denies. But from the standpoint of our ordinary practice, there’s nothing wrong with that, taken all by itself. Suppose that someone contends that you don’t know anything about George W. Bush’s gender. And suppose that this person invites you to consider the hypothesis that Bush is a woman successfully engaged in a massive deception. The *right* response, given how things are, is this. “Don’t be ridiculous; Bush is a man.” I grant that this response might not be dialectically effective in convincing your interlocutor, and so it may not be the best thing to say if you hope to bring about conviction. But in saying that this is a *right* response, I only mean that it is an appropriate response that you are fully entitled to make: the mere assertion that one doesn’t know something about a certain domain doesn’t preclude one from making use of things that one knows in that domain. The same goes for *considering the possibility* that one doesn’t know anything in that domain. (Sometimes one’s serious and sincere consideration of a

suggestion amounts to little more than summarily rejecting it.) So something more than mere assertion or consideration of the skeptical conclusion will be needed in order to explain why I can't appeal to considerations about the world in response to the skeptic's use of the evil demon hypothesis. In particular, there will need to be something to be said in favor of the skeptic's crucial claims or key moves.

The Adequacy of our Evidence

Here's an argument that is sometimes offered in favor of the second premise of the Argument from Ignorance: you don't know (or reasonably or justifiably believe) that you aren't being deceived by an evil demon, because if you were being deceived by an evil demon you would believe (incorrectly) that you aren't. This thought is codified by Nozick in his Sensitivity condition on knowledge: to know *p*, it must be the case that if *p* were false, you wouldn't believe *p* (ref.). But the Sensitivity condition is incorrect. A number of counterexamples to it have been offered over the years (ref's). Here's a simple one, inspired by Timothy Williamson [ref]. Suppose that you are looking at a twenty-foot pole. Your perceptual experience of the pole enables you to know that it is more than ten feet tall. Now let's stipulate that the situation is such that if the pole were not more than ten feet tall, it would be nine and a half feet tall. (Perhaps the pole comes in two sections, the first of which is nine and one half feet tall. The workmen had to decide whether to erect only the first section or both, and they could easily have gone either way.) And let's stipulate that you aren't very good at estimating heights visually, so that if the pole were nine and a half feet tall you would believe that it is more than ten feet tall. So this is a case in which, looking at the twenty foot pole, you know that it is at least ten feet tall, but if the pole were not at least ten feet tall, you would still believe that it was. So your belief is not Sensitive, but you have knowledge nonetheless. It should be obvious that the same point would apply to reasonable and justified belief, as well as to the notion of having good or adequate reason to believe.⁷

Even if Sensitivity isn't required, though, we might nonetheless lack adequate evidence that we aren't being deceived by an evil demon. Here's another proposal that we sometimes hear: we aren't entitled to believe that we are not being deceived by an evil demon because our evidence is neutral on the question. And here's an argument that is sometimes offered in support of that suggestion: if one was being deceived by an evil demon, one would have the very same evidence that one has now.

This argument makes use of a principle that looks something like this: if I would have exactly the same evidence that I have now even if *p* were not the case, then my current evidence is neutral on the question of whether *p*. That principle is incorrect. It is surely possible for the following to happen: I believe that *p* on the basis of fallible (defeasible) evidence that *p*, and the situation is such that if *p* were false, I would still have just that evidence (because the defeating condition's obtaining would be unknown to me). Any such situation is a counterexample to the principle.

However, there is a much more fundamental point to be made here. It is not true that my evidence right now is the same as the evidence I would have if I were being deceived by an evil demon. Here is a piece of evidence that I currently possess: there are no evil demons. That claim would not be true if I were being deceived by an evil demon. So my evidence would be different if I were being deceived by an evil demon. I thus currently have evidence that I would lack under those circumstances, and this evidence is not neutral on the question.

Of course, if I were being deceived by an evil demon I would *incorrectly believe* that there are no evil demons. So it would *seem to me* that I had exactly the evidence that I have now. But that fact does not show that I now lack that evidence. Here's an example that can help make the point. Suppose that when I am drunk, I slur my words while believing – incorrigibly so – that I am not doing so. (I incorrectly believe that I hold my liquor so well that I speak impeccably even when drunk, and when drunk I unconsciously misinterpret everything in a way that accords with this belief.) So when I am drunk, it seems to me that I have just the evidence that I now have that I am not slurring my words. Right now, however, I have evidence that I am not slurring my words. So I have evidence right now that I would lack if I were drunk and slurring my words, despite the fact that under those conditions I would believe that I had that evidence. No one thinks that the fact – that, were I drunk, it would seem to me that I have the evidence that I now have that I'm not slurring my words – forces us to conclude that I don't now know that I'm not slurring my words. Why think any differently about the evil demon case?

In order to establish the charge that my evidence is neutral on the question of whether I'm being deceived by an evil demon, you have to somehow place limits on what counts as my evidence. You have to preclude me from including amongst my evidence such facts about the world as that there are no evil demons. But how is that limitation to be accomplished? This is just another way of putting the question with which I started. Are there any commitments or principles from our ordinary practice that put considerations about the world out of play when we encounter the skeptical argument?

It might be suggested that the required restriction is generated by a feature of our ordinary practices of epistemic evaluation: what we count as a person's evidence that *p* is relative, it might be said, to the alternatives to *p* that are under consideration. So consider a hypothesis *H* which implies that *S* doesn't know that *p*. It might be proposed that when one considers such a hypothesis, one thereby restricts what counts as *S*'s body of evidence to just the evidence *S* would have whether or not *H* is true.⁸ This proposal could be given an invariantist cast, according to which it reveals an underlying fact about what the evidence actually *is* even when we are not considering the hypothesis *H*. Or it could be given a contextualist cast, according to which the truth-conditions of evidence attributions are sensitive to the possibilities under consideration by the attributor.⁹ Either way, though, the proposal has the consequence that we could not acceptably appeal to or rely upon considerations about the world in order to counter the evil demon hypothesis.

This proposal is incorrect. Consider again the hypothesis that George W. Bush is a woman successfully engaged in a massive deception. This hypothesis implies that you do not know that Bush is a man. Does considering it limit your evidence in the suggested way? Obviously not. Bush was married to a woman, Laura Bush, in Texas – a state which does not perform same-sex marriages. That’s a piece of evidence which you would not have if the hypothesis were true. But it’s a piece of evidence you have that Bush is a man, and you have it even now that we are considering this hypothesis. Of course, it would not be dialectically effective to present this evidence to someone who was genuinely inclined to believe the hypothesis. But the hypothesis is patently false, and so considering it does not lead us to restrict our evidence in this way. We *would* appropriately restrict our evidence in this way if there were some good reason in favor of this hypothesis. But the parallel point regarding the evil demon hypothesis won’t help the skeptical argument, because that hypothesis doesn’t have any reason in its favor.

Epistemologists sometimes suggest that we value something which might be called “reflectively accessible evidence,” and that the desirability of such evidence underwrites special “internalist” conceptions of knowledge, justification, and other terms of epistemic approbation. In the intended sense of the phrase, “reflectively accessible evidence” is not just evidence that is available to you when you pause to reflect on the question, “What reason do I have to hold this belief?” It is not just something that you are in a position to cite. Rather, it is a justification that you possess, and that you can identify as such, when you bracket all of your convictions about the world and utilize only the resources of introspection and a priori reflection. Understood in that way, it is pretty plausible that we lack “reflectively accessible evidence” for the belief that we are not being deceived by an evil demon. Might this consideration reasonably move us not to rely on our ordinary conviction that there are no evil demons?

I don’t think it can do the trick. Of course, we can define whatever special conception of evidence or “justification” we like. The question of interest for my project here is whether the special notion, so defined, is relevantly connected with the values that appear in our ordinary judgments about what people know, ought to believe, and the like. In this case, I think the answer is pretty clearly, “No.” I do not have a reflectively accessible justification – a good reason that is provided purely by the resources of introspection and a priori reflection and that is recognizable as such using only those resources – for thinking that George W. Bush is not a woman successfully engaged in a massive deception. But for all that, my belief on this matter is in excellent epistemic shape. Moreover, there aren’t any examples that appropriately connect our lack of a “reflectively accessible justification” with the evaluations that we ordinarily make. Some epistemologists have argued that examples of “chicken sexers” and the like – people who have capacities to make reliable or counterfactually secure judgments about a certain subject matter while having no reasons available to their conscious reflection and no correct beliefs about the factors that guide these judgments – show that we especially value “reflectively accessible justifications.”¹⁰ But these examples don’t show any such thing. If we judge such subjects to be in a sub-

optimal epistemic position, that can be explained in terms of the fact that they lack reflectively accessible reasons, in the *ordinary* sense of that phrase: they can't cite any reasons in favor of their judgments. Our responses to such examples thus don't involve a demand for "reflectively accessible reasons" in the special sense defined above.¹¹ The demand for a "reflectively accessible justification", understood in the special way defined above, is not something we are committed to in our ordinary practice.

Epistemic Closure and Epistemic Priority

It is sometimes argued that once you accept the first premise of the Argument from Ignorance, you are sunk: you are forced to accept the second as well. (This is a way of summarizing the burden of Stroud's argument in SPS, pp. [].) If that's right, then something about the first premise, or about the considerations lying behind it, enforces the limitation that I have been highlighting.

On many accounts, the first premise is underwritten by the *principle of epistemic closure*: roughly, if you know that p and know that p implies not-q, then you know that not-q. Similar principles might be proposed for justified belief and for having reason to believe. For the sake of argument, let's grant that these principles are correct. Might they hold the key to the question I've been pressing?

I have to admit that it has always puzzled me that so much weight is placed on the closure principle in discussions of skepticism. As Stroud pointed out long ago, the skeptical argument would seem to work just as well with possibilities that are *compatible* with the truth of many of my beliefs about the world [ref]. Consider, for instance, the hypothesis that I am an evil demon's plaything: sometimes it gives me false beliefs and sometimes true ones, but it does so in a way that ensures that I am awfully unreliable about what's going on around me. (We might imagine that the evil demon supplements these deceptions by sometimes making me disembodied, so that I'm not even reliable about such matters as whether I have hands.) Now consider my belief that there is a piece of paper in front of me. That there is a piece of paper in front of me does *not* imply that I am not an evil demon's plaything. So the closure principle does not even apply here. But still, now that I've thought about this possibility, the first premise of the Argument from Ignorance seems just as plausible regarding it. And it seems that a skeptical argument employing this possibility would have just as much plausibility as one employing the possibility that all my beliefs about the world are false because I am *always* being deceived by an evil demon. These considerations strongly suggest to me that the emphasis on the closure principle is misplaced in discussions of skepticism: something else must be doing the work.¹²

Still, let's consider the closure principle and the possibility that one is a victim of an evil demon who *always* deceives (so that all of one's beliefs about the world are false). Does the closure principle, in conjunction with this possibility, force us to accept the second premise of the Argument from Ignorance?

No. The relevant instance of the closure principle states a necessary condition on knowing certain things: it says that in order to know those things, you also have to know that you are not being deceived in this way by an evil demon. That's no problem, at least not so far. Stroud has argued that a premise along the lines of premise 1 becomes unsatisfiable when applied to the question of whether one knows that the skeptic's hypotheses don't obtain (SPS, ch. 1). But that's not so, if premise 1 is underwritten by the closure principle. When we apply the relevant instance of the closure principle to our knowledge that we are not so deceived, we get this (ignoring the clause concerning knowledge of the relevant implication):

I don't know that I am not being deceived by an evil demon unless
I know that I am not being deceived by an evil demon.¹³

And that, of course, just states the truism that you can't know something unless you know it. That's no problem.

In Stroud's version of the skeptical argument, the closure principle is replaced by this: if you know that *q* is incompatible with your *knowing* that *p*, then in order to know that *p*, you must know that not-*q* (SPS, []). (This principle is designed to allow for the fact that the skeptical hypothesis apparently need not be incompatible with the truth of all of the relevant beliefs.) This principle likewise fails to yield an unsatisfiable requirement when applied to the case of one's knowledge that one is not the victim of an evil demon. The problem in both cases is simple: you can't get a vicious circle or regress out of a mere necessary condition; *p* is trivially and innocuously a necessary condition of itself. That's why the closure principle (or Stroud's variant) alone won't get you the second premise of the Argument from Ignorance.¹⁴

Here's a thought that might do the trick. It seems pretty plausible that I won't know that I'm not being deceived by an evil demon unless I also know a bunch of other things about the world, such as that there aren't any evil demons (and whatever else I need to know in order to know that). Conjoin that thought with the first premise of the Argument from Ignorance (perhaps underwritten by the closure principle). Then you get this.

- i. In order to know other things about the world, I need to know that I am not being deceived by an evil demon. (premise 1)
- ii. In order to know that I am not being deceived by an evil demon, I need to know other things about the world (e.g., that there aren't any evil demons). (the additional thought)

Don't (i) and (ii) create a vicious circle?

Well, no. All that they amount to is that to know some one thing I've got to know a bunch of other stuff as well, and that in order to know that other stuff I've got to know the first thing. And there's nothing problematic about that. We can perfectly well make sense of a situation in which you can't know that *p* unless you know that *q* and also can't know that *q* unless you know that *p*. (The same goes for justified belief and for having reason to believe.) So no problem arises even when we conjoin the closure principle with the thought that our knowing that we are not being deceived by an evil demon *depends upon* our knowing other things about the world.

A problem *would* arise here if (i), (ii), or both involved an *epistemic priority requirement*: that is, a requirement to the effect that certain things must be known antecedently to or independently of other things. For if either (i) or (ii) involved a priority requirement, then (assuming that the relevant priority relation is transitive) the result would be that in order to know that I am not being deceived by an evil demon, I must antecedently/independently know that I am not being deceived by an evil demon (or, alternatively, that in order to know those other things about the world, I must antecedently/independently know those other things about the world). And that *is* an unsatisfiable requirement.

Now it might seem that I have already committed myself to a priority reading of (ii). I have appealed to the consideration that there are no evil demons as a ground for dismissing the possible hypothesis that one is being deceived by an evil demon. This could easily be understood as the suggestion that the consideration that there are no evil demons is the ground for an inference through which I derived my knowledge or justified belief that I am not being deceived by an evil demon. So understood, a problem regarding epistemic priority would arise. I couldn't then also grant that in order to know or have justified belief that there are no evil demons I would already have to know or justifiably believe that I am not being deceived by an evil demon.

This problem does not arise, however. To appeal to a consideration as a ground for *dismissing a possible hypothesis* is not necessarily to say that the former consideration serves as the ground for an inference through which I have derived my knowledge that the hypothesis is incorrect. To provide a ground for dismissing a possibility is to provide a consideration that warrants treating that possibility in a certain way in the course of the deliberation. This is different from inferring from some evidence or premise that the possibility does not obtain, and it is also different from reporting a prior inference upon which one bases one's belief that the possibility does not obtain. For instance, a gay friend of mine recently had an acquaintance exclaim with surprise, "But you're straight!" "Of course I'm not," she replied. "I've been dating the same woman for five years." In so dismissing the suggestion, she was not inferring from that consideration that she is not straight, nor was she reporting an inference through which she had arrived at that knowledge. Dismissing a possibility on a particular ground and inferring or reporting an inference are thus different things. Moreover, one can appropriately provide a ground for dismissing a possible hypothesis without thereby creating or tracking an epistemic priority relation amongst the beliefs in question or amongst the propositions believed. For instance, I am completely confident that my wife loves me. I don't have any doubt or reason for doubt or suspicion. Now I open a fortune cookie that says, "Your wife doesn't love you." I may quite appropriately dismiss the possibility as ridiculous. She shows me she loves me in hundreds of ways every day, and she wouldn't deceive me about such matters. Suppose now that as I leave the restaurant, I see a man with a sign that reads, "Your wife lies when she says she loves you!" I may quite appropriately dismiss that thought as ridiculous too. She loves me too much to deceive me about such important matters. There's obviously nothing problematic about my offering *both* of these responses. If appropriately providing a ground for

dismissing a possible hypothesis always created or tracked an epistemic priority relation, then this example would involve an objectionable circle. But it doesn't. So one can appropriately provide a ground for dismissing a possible hypothesis without thereby creating or tracking an epistemic priority relation.¹⁵

Our situation, then, is this. I have granted that my knowing that I am not being deceived by an evil demon requires my knowing other things about the world, and I have appealed to some of those other things as a ground for dismissing the suggestion that I am being deceived by an evil demon. But none of that requires me to regard my knowledge of those other things about the world as epistemically prior to my knowledge that I am not being deceived by an evil demon.¹⁶

What about a priority relation running in the other direction? Suppose that premise (1) of the Argument from Ignorance stated an *epistemic priority condition* to the effect that in order to know any particular thing about the world, one must antecedently know that one is not being deceived by an evil demon. From that premise there would be a quick path to the skeptical argument's second premise. For then, applying the first premise to the case of one's knowledge that one is not being deceived by an evil demon, we would get the result that one cannot know that one is not being deceived by an evil demon unless one antecedently knows that one is not being deceived by an evil demon. That's an unsatisfiable requirement. And this result would buttress the charge that it is question-begging to reply to the second premise of the argument by appealing to considerations about the world, for one would *antecedently* have to know that one is not being deceived by an evil demon in order to make use of those considerations.

Is there any reason to accept the first premise, read as an epistemic priority requirement?

There would be if a *priority* version of the closure principle (or of Stroud's variant of it) were correct. For then the priority version of premise (1) could be *derived* from that more general priority principle.

Here's what the general priority version of the closure principle would say:

If you know that p and know that p implies not-q, then you must antecedently know that not-q.

But this principle is clearly false. For one thing, it would preclude the very possibility of arriving at knowledge through deductive inference. Even more radically, however, it would straightforwardly make all knowledge impossible. Here's a simple example which provides a template applicable to any proposition you like. Suppose that you are looking straight at a cow in normal conditions. You recognize that it's being a cow implies that it is not a chicken. So according to this principle, to know that it is a cow, you must antecedently know that it is not a chicken. Now consider your knowledge that it is not a chicken. You recognize that its not being a chicken implies that it is not a chicken. So according to this principle, to know that it is not a chicken, you must antecedently know that it is not a chicken. But that's impossible. So you can't know it's not a

chicken. And so you can't know it's a cow. This argument obviously generalizes. And it applies to Stroud's variant of the closure principle as well. If our aim is to show that the priority version of premise (1) is a commitment of our ordinary practice, this isn't the way to do it.[refs: Pryor, Byrne, me]¹⁷

Begging the Question, second pass: trying to ground a priority requirement

Here's a more promising strategy for deriving the needed priority requirement. If I *were* being deceived by an evil demon, then my reliability, competence, or authority would be compromised regarding a wide range of matters, including such things as whether there are evil demons. It might consequently be thought that in order to be entitled to believe anything about that range of matters, I would need to have adequate grounds, not in that range, for believing that I am not being deceived by an evil demon. And if that's so, then the problem would be this: in appealing to considerations about the world in response to the evil demon hypothesis, I wouldn't be addressing the right question; I wouldn't be showing that I have *the right kind* of grounds for believing that I'm not being deceived, grounds of a kind that I would have to have in order to be entitled to believe anything about the world at all.¹⁸

What's going on here can be usefully characterized in terms of the notion of *independent grounds*. Suppose that a certain hypothesis H is such that if it obtained, then one's reliability, competence, or authority would be compromised across a certain range of propositions D. Say that a ground is *independent* if it is not in this domain. We can then formulate a requirement schema as follows:

If H is a hypothesis whose truth would undermine my reliability, competence, or authority regarding some domain D, then I ought not believe anything in D unless I have adequate independent grounds for believing that not-H.

The upshot, if such a requirement is correct, is an epistemic priority structure between the belief that not-H and considerations in the domain: one cannot acceptably believe anything in the domain without having adequate grounds from outside the domain for believing not-H. (I will consequently use the term "excluded domain" to refer to this domain.)

Of course, being deceived by an evil demon would compromise one's reliability, competence, or authority regarding a wide range of propositions. So, applying this requirement to this case, we get this:

For any proposition about which I would be unreliable, incompetent, or lack authority if I were being deceived by an evil demon, I ought not believe it unless I have adequate independent grounds – adequate grounds from outside this domain – for believing that I am not being so deceived.

Now it's at least arguable that I lack such grounds. If so, then (given this requirement) I shouldn't believe anything in this domain. So I wouldn't possess anything, either in the excluded domain or outside it, that would provide an adequate ground for believing that I am not being deceived by an evil demon. And in that case, I ought not believe that I am not so deceived. Could this explain why we can't reject the evil demon hypothesis on the basis of considerations about the world around us?

I don't think so. To see why, let's start by considering this question: what might be meant in this context by the phrase "adequate grounds for believing"?

Here's one thing that might be meant: grounds sufficient for a fully rational, fully reflective deliberator to arrive, through a fully explicit and conscious course of reasoning, at the conclusion that she is not being deceived by an evil demon. But if that's what's meant, then the requirement is surely incorrect. For then we would be understanding the requirement schema like this:

If H is a hypothesis whose truth would undermine my reliability, competence, or authority regarding some domain D, then I ought not believe anything in D unless I can acceptably arrive at a belief that not-H through an explicit, conscious course of reasoning from adequate independent grounds.

And the schema, understood in this way, yields the wrong results in perfectly ordinary sorts of cases.

Consider again the hypothesis that George W. Bush is a woman successfully engaged in a massive deception. Do I have grounds from outside the domain that enable me to arrive through an acceptable course of explicit, conscious reasoning at the belief that George W. Bush is not a woman successfully engaged in a massive deception? I don't think so. But of course I am entitled both to dismiss this hypothesis and to believe that Bush is a man.

Here is a consideration that seems to tell pretty directly against that hypothesis: While deceptions of this sort have taken place, they are extremely unusual and extraordinarily hard to pull off. That's true. But taken by itself, it isn't a sufficient reason for concluding that the possibility doesn't obtain. Suppose that you are given the following information: Events of type X are extremely unusual and require very special sorts of circumstances. You are given *no other relevant information*. Now you are asked, "So, is an event of type X taking place right now?" It's a pretty safe bet that such an event isn't taking place, and you are in a position to say, "Probably not." But you shouldn't conclude from such information that an event of type X isn't taking place now. To reach that conclusion, you need more background information. What more would you need? That's not clear. But here's something that would do the trick: enough background information about the circumstances for you to be in a position to conclude that nothing tells in favor of the hypothesis. In the example we are considering, however, that claim is itself something that is in the excluded domain. And if we consider only the information that you possess outside the excluded domain, it isn't enough to enable you to reach that claim.

It might be objected, however, that there is lots and lots of evidence from outside the excluded domain that *George W. Bush is a man*. For instance, Bush looks like a man. Bush sounds like a man. Bush walks like a man. Members of the Congress, White House staff, and Bush family all refer to Bush as "he." It might be suggested that this evidence provides adequate ground for believing that George W. Bush is a man. And if this is so, then – it might be suggested – we can acceptably reason from grounds outside the excluded domain to the belief that George W. Bush is not a woman successfully engaged in a massive deception;

after all, this conclusion is a deductive consequence of George W. Bush's being a man, and we have lots of evidence that that's exactly what Bush is.

The objection imagines a two-step inference that runs like this:

1. Bush looks like a man, etc.
2. So, Bush is a man.
3. So, Bush is not a woman successfully engaged in a massive deception.

Now something is pretty clearly wrong with this line of reasoning. To see this, just imagine someone who does not yet have any attitude – not belief, inclination to believe, doubt, inclination to doubt, or even suspension of judgment – or any reason to yet adopt any attitude regarding the proposition that Bush is a woman successfully engaged in a massive deception. They just don't have any attitude or any information either way. Now they want to find out how things are: whether Bush is a woman successfully engaged in a massive deception. They start by observing that Bush looks like a man, talks like a man, walks like a man, and is widely referred to as "he". They recognize, of course, that if Bush were a woman successfully engaged in a massive deception, then these considerations would be evidentially worthless. Suppose, now, that recognizing all this, they go on to infer that Bush is a man. There will be a moment, just before they form the belief that Bush is a man, at which their deliberative position would be fully represented like this: "I recognize that if Bush were a woman successfully engaging in a massive deception, that would prevent these considerations from telling in favor of the truth of the claim that Bush is a man, and as of yet I have no opinion whatsoever as to whether Bush is a woman successfully engaging in a massive deception, but still, I conclude from these considerations that Bush is a man." This position is pretty clearly unsatisfactory: someone proceeding in this way would be open to epistemic criticism.

On my view, what goes wrong here has to do with the commitments that one undertakes when one makes the inference from the evidence to the conclusion that Bush is a man. When one consciously and explicitly infers from *p* to *q*, one thereby commits oneself to taking *p* to tell in favor of the truth of *q* [Leite 2008 and forthcoming2]. In many cases (including the one at hand), what something tells in favor of will depend upon other conditions in the world. If one recognizes that *W* is a way in which *p* could fail to tell in favor of the truth of *q*, then on pain of inconsistency, one is committed to taking *W* not to obtain if one makes such an inference. So someone who makes the inference about Bush in the circumstances imagined in the example is thereby already committed to taking Bush not to be a woman successfully engaged in a massive deception – and is so committed as a condition of making the inference acceptably at all. If they lack adequate grounds in support of this commitment, however, then their position regarding their conclusion will be open to epistemic criticism [Leite forthcoming2]. And that, I would suggest, is what is going wrong in the example. In order to reach the first step of the inference, the person already needs a well-supported belief that Bush is not a woman engaged in a massively successful deception, and so the second step of the inference would be otiose. Moreover, in the case we are imagining the person lacks a well-supported belief to that effect, so the first step of the inference cannot be acceptably made. And analogous problems arise if we consider the

one-step inference: Bush looks like a man (etc.); so, Bush is not a woman successfully engaged in a massive deception.

I consequently conclude that in order to believe such things as that Bush is a man (a proposition in the excluded domain), one does not need to have independent evidence – evidence from outside the excluded domain – that suffices for one to come to believe through an acceptable course of explicit, conscious reasoning that Bush is not a woman successfully engaged in a massive deception. This version of the priority requirement-schema should consequently be rejected.

Let's pursue the argument one step further. Someone drawn to the sort of requirement we've been considering might say that it doesn't matter if the person couldn't explicitly reason in an acceptable way from the independent grounds in question to the conclusion that the relevant hypothesis is incorrect; all that's required is that they *have* those independent grounds (and perhaps that their belief-forming systems be appropriately responsive to them). So, for instance, it might be urged that Bush's looking like a man constitutes independent grounds that *do* decisively tell against George W. Bush's being a woman successfully engaged in a massive deception, even if we can't acceptably explicitly reason from those grounds to that conclusion. It might be suggested (i) that if that weren't so we wouldn't be entitled to believe such things as that George W. Bush is a man, and (ii) that it's in part because it is so that we can reject this hypothesis by appealing to such considerations as that George W. Bush is a man. And it might be suggested that this example differs from the evil demon possibility on precisely on this score: we do not possess independent grounds that decisively tell against the hypothesis that we are being deceived by an evil demon.

The suggestion, then, is that we should interpret the schema like this:
 If H is a hypothesis whose truth would undermine my reliability, competence, or authority regarding some domain D, then I ought not believe anything in D unless I have independent grounds which in fact decisively tell in favor of not-H.

On this interpretation, the schema only requires that I have the relevant grounds, even if I can't acceptably arrive at a belief that the hypothesis is incorrect through an explicit, conscious course of reasoning from those grounds. Why does this difference matter? Because it seems at least possible that one could possess grounds which in fact decisively tell in favor of the truth of a proposition even though one could not acceptably arrive at a belief in the truth of that proposition through a course of conscious, explicit reasoning from those grounds.

But is even this requirement something we are committed to by our ordinary practice? I think not. Consider the following hypothesis:

Children of at least one brunette parent are massively bad at evaluating and responding to evidence, both in conscious deliberation and in non-deliberative belief formation, and they also suffer from widespread and significant deficits and distortions in sensory processing (compensated for by "infilling" and confabulation).

This hypothesis is clearly false. We would all dismiss it out of hand. And if you are like me in having a brunette parent, then in dismissing it you would be relying upon considerations about which you would lack reliability, competence, or authority if the hypothesis were true. Consequently, if you are like me there is -- among the things we believe -- no independent evidence (in the sense I defined above) that decisively tells against this hypothesis. Moreover, suppose we grant that one's evidence can include not just things one believes, but also one's mental states, especially sensory experiences. It might be said that as things are my actual sensory experiences do provide decisive evidence against this hypothesis. Still, if the hypothesis were true, I would lack important sensory competences. So even my sensory experience (now) is in the excluded domain; it does not count as "independent" relative to this hypothesis. I consequently lack even sensory evidence from outside the excluded domain that decisively tells against the truth of this hypothesis. But I am nonetheless entitled to reject this hypothesis, and in support of doing so I would appeal to considerations in the domain excluded by the requirement in question (starting with the fact that there is no reason whatsoever in this hypothesis's favor!).

I think that this should be the end of the matter: we should reject the requirement even under this minimal interpretation. It might be suggested, however, that this example doesn't really tell against the requirement because it fudges by putting one's current sensory evidence into the excluded domain. Maybe the example actually *supports* the requirement: maybe we *do* have sensory evidence that in fact decisively tells against the hypothesis, and that's why we are entitled to reject it. Suppose, for the sake of argument, that we grant this response. It doesn't get us very far in defending the skeptic. For the same thing could now be said in response to the evil demon hypothesis: we *do* have sensory evidence that in fact decisively tells against this hypothesis (for instance, one might say that given how the world works, our sensory evidence *does* reliably indicate the truth of the proposition that there are no evil demons), and that's why we are entitled to reject it.

The situation, then, is this. Even if I'm wrong that this version of the priority requirement is incorrect, there's no hope of utilizing it to get a skeptical argument going unless we can find some reason to distinguish the evil demon hypothesis from the children-of-brunettes hypothesis in the relevant respect. I do not think this can be done. It might be suggested that the necessary work is done by the addition of this thought: if you *were* being deceived by an evil demon, then your sensory experience wouldn't be any evidence at all that there are no evil demons, since your sensory experience wouldn't be appropriately connected with the truth about the matter. But for this thought to help the skeptic, we also need *this* thought: because of that fact, your sensory experience *now* is not evidence that there are no evil demons. And that thought depends on the following principle:

If E would not tell against the truth of p if p were true, then E does not tell against the truth of p.

The trouble here is that this principle cuts against the suggested response to the children-of-brunettes hypothesis as well. *Of course* we can acceptably reject out

of hand the hypothesis that children of brunettes suffer from widespread defects in sensory processing and are massively bad at dealing with evidence. If you say that that's because as things are my sensory evidence decisively tells against that hypothesis, then you can't accept this principle. So I don't see any hope of getting a reasonable skeptical argument going in this way.

In sum: I've been considering a line of thought that tries to get a priority requirement out of the recognition that if we were being deceived by an evil demon, we wouldn't be reliable, competent, or authoritative over a wide range of matters. I've argued that this line of thought won't generate a priority requirement that does the work required for the skeptical argument. It won't succeed in moving us from our ordinary position to a position in which we deem it inappropriate to rely upon considerations about the world when determining what we know or what we ought to believe.

The epistemologist's special question?

A number of philosophers, most prominently Barry Stroud and Laurence Bonjour, have suggested that what moves us from our ordinary position is a distinctive question or project that sets the stage, as it were, upon which the skeptical reasoning takes place. Stroud suggests that the distinctive project is one of investigating "human knowledge in general", and that the project is initiated by the question, "How does anyone know anything at all about the world around them?" [ref] Bonjour suggests that the fundamental question instead concerns our reasons: "Do we have good reasons for believing anything at all about the world around us?" [ref] Both suggest that once the question, understood in the intended way, is taken up as a matter for investigation, we are precluded from making use of any considerations about the world in the course of answering it. So the suggestion is that it is the investigation itself that moves us from our ordinary position to a position in which considerations about the world around us are off the table. And once that shift has taken place, we are no longer in a position to reply to the skeptical argument in the way I have been urging.

For this proposal to succeed, it has to be straightforward to get the special investigation going from within our ordinary standpoint. But if we start off in our ordinary position and are clearheaded about what that involves, it really isn't so easy to bring about this shift. To make the point, I want to focus on Bonjour's version of the project.¹⁹ A brief sketch of my concern should suffice to bring out the main issue.

To begin, consider that in the relevant sense of "reason", anything that counts as a reason for a belief about the world will be good evidence: it will tell in favor of the truth of the belief in question. But in a great many cases, what tells in favor of what depends upon facts about the world, including facts about what else is the case, how things tend to work, what is reliably correlated with or caused by what, and the like. For instance, and to take a particularly pertinent example, if you were being deceived by an evil demon, then your sensory experiences wouldn't be evidence *at all* about what is the case in the world around you (though you would wrongly take them to be). I take it that this fact about evidential support relations is something that we are committed to in our ordinary

position: it is a commitment that appears in our ordinary practices of reasoning and of defending and evaluating our own and other people's beliefs.

Now the question Bonjour proposes is perfectly intelligible. But notice that once we have in clear view the point I just made about epistemic reasons, we can see that the question – understood as Bonjour intends – is unanswerable. It is unanswerable because once one has precluded oneself from appealing to considerations about the world, one lacks the resources one needs in order to determine what, if any, reasons one has for beliefs about the world around us.²⁰ This result is purely a product of the interplay between the structure of the question and the fact about epistemic reasons I just highlighted. It has nothing whatsoever to do with whether we have any reasons for believing anything about the world around us. So there is no good reason to go in for the project this question initiates. Look at it this way. We can see that the project must fail and that its failure provides no answer to the question that got it going. We don't have to be able to answer that question, understood in the intended way, in order to have good reasons for believing things about the world around us. And from our ordinary position, making use of things we know about the world, we can perfectly well determine both whether we have good reasons for believing particular things about the world and what those reasons are. From that vantage point, then, going in for the project Bonjour describes just looks like a waste of time and effort, like trying to determine by sight how many objects are in a completely dark room. Not every question deserves a serious attempt at an answer.

For this reason, the suggested question does not move me from my ordinary position. When I consider it, I don't find myself thinking that I have found something, accessible from my ordinary position, that precludes me from appealing to considerations about the world in response to skeptical arguments. Instead, I conclude that I have found a proposed research project that there is no good reason to go in for.²¹

Of course, if you weren't considering Bonjour's question from that standpoint, then you won't feel the force of this reply. But then you also won't be engaging with the question that I've meant to be pursuing today.

Postlude

I want to close by briefly reflecting on what I have done today. I have delayed such reflections until now because I wanted us to keep our feet squarely planted in ordinary life, and it's hard to do that when we begin by thinking about how the project is to be placed within a larger philosophical context.

My arguments (proceeding as they did from within our ordinary standpoint) relied heavily upon claims about the world around us and about what we currently know. In this regard, they may have seemed reminiscent of Moore's famous attempt to refute skepticism [ref], and they may have struck some of you as similarly dogmatic. However, unlike Moore, I was looking at an earlier stage in the skeptical argument and asking how the crucial first steps of that argument are supposed to get going. At that stage in the dialectic all of our ordinary commitments are still appropriately in play, and the question is whether some of

our them will put others out of play. For the skeptic who attempts to work from within our ordinary position, there is no alternative at this stage than to appeal to our commitments regarding particular examples, and so that's where I looked as well. I did not argue that skepticism *cannot* arise out of those materials. Rather, I pointed out that so far no reason has emerged to think that it does.

Many contemporary “Moorean” responses to external world skepticism – for instance, Jim Pryor’s – depend upon a particular theory of empirical justification. According to this theory, there is *immediate empirical justification*: one’s sensory experiences can give one justification to believe something in a way that does not constitutively depend upon any of one’s justified beliefs [ref Pryor]. The arguments that I offered are compatible with this theory.²² But they do not require it, and they are equally compatible with its rejection. This point highlights an important aspect of my approach. From my point of view it is a mistake to think that full-blown epistemological theorizing is needed before we can engage in the kind of response to skepticism that I have considered. A negative response to the skeptic who works “from within” would be one that we are entitled to make by the norms and commitments that are involved in our ordinary practices. What an epistemological theory would do is clarify, articulate, and systematize the relevant features of our ordinary practice (perhaps also with an eye to improving it in various ways, if we come to think that is best). The arguments that I offered here would thus provide *constraints* for such theorizing. Of course, it’s possible that arguments might be offered for discounting some or all the considerations I’ve put forward. We would have to take up such arguments as they arise.

To put it programmatically:

I want to urge the *primacy of our ordinary epistemic practice (including our best scientific practice)*. Start from where we are in the midst of our ordinary lives, and no reasonable skeptical argument will get going. Philosophical reflection may give us theoretical insight into why that is so. But that it *is* so is a pretheoretical constraint that our best theorizing must either respect or somehow explain away.

¹ A further, very important question is what significance this result would have for the question of whether skepticism is true. I will not broach that further question here.

² Others who I take to be heading in the same direction: Ernie Sosa, Alex Byrne, Conee/Feldman, Matthias Steup, and Penelope Maddy, and in their own ways G. E. Moore, J. L. Austin and W.V. O. Quine beat us to the punch line long ago

³ ref

⁴ This presentation will idealize away from possible concerns about degrees of belief or credence and will work with an on/off notion of belief.

⁵ For a full-dress discussion of this form of response to the dream-argument for skepticism, see my [forthcoming1].

⁶ I’m grateful to Katy Abramson for highlighting this way of seeing my project.

⁷ It might be complained that this counterexample doesn't properly parallel the evil demon hypothesis, since in the counterexample one's perceptual experiences in the two cases would differ radically when described "from the inside".

However, there are counterexamples to the sensitivity requirement that do not require the cases to differ in this way. Suppose that a play opens with a scene in which the lead walks by what appears to be a closed book sitting on a table. The stage manager for the production is extremely conscientious: he wouldn't allow the actor to go on unless the scene is properly set. As the actor knows, the stage manager has a very strong policy of using real objects rather than fake props. And of course books are in plentiful supply, so there's no need to resort to a painted wooden or papier mache prop. The actor walks onto the stage, looks at the book on the table and comes to believe that there is a book on the table. He thereby comes to know that there is a book on the table. If there weren't a book on the table, however, there would be a fake prop that looks like a book. In that case he would believe that there is a book on the table. So his belief that there is a book on the table isn't Sensitive. But still, in the actual case in which books are plentiful and there is a book on the table, he knows that there is a book on the table.

⁸ This is the core suggestion behind Ram Neta's rule "R!" in "S knows that p," NOÛS 36:4, 2002, 663–681, at p. 674. Neta's formulation is considerably more complicated, for reasons arising from the details of his theory of evidence.

⁹ This is more or less Neta's proposal.

¹⁰ E.g., Duncan Prichard, pp. 174-5 and *passim*. It should be noted that to be relevant at all, the examples must be understood in such a way that the person can't even provide track-record considerations in favor of believing as she does.

¹¹ What *would* show this is an example of the following sort: though a person has reasons in the form of considerations about the world that she can cite in defense of her belief, using *only* the resources of introspection and a priori reflection she cannot come up with good reasons for believing as she does—and we judge that her belief is consequently in poor epistemic shape. But again, I don't think that there are any such examples. I can cite reasons in favor of my belief that George W. Bush is not a woman successfully engaged in a massive deception, but I cannot provide good reasons for believing as I do if I am limited only to the resources of introspection and a priori reflection. Nonetheless, my belief is just fine.

¹² There are other significant questions to be raised about whether the closure principle can play the role that it is widely held to play in skeptical arguments. For extensive discussion of this issue, see part [] of my 2004. As I argue there, the closure principle would in fact be entirely idle in the skeptical argument.

¹³ Let $p = \sim e.d.$ (the evil demon hypothesis is not the case). Let $q = e.d.$ (I am being deceived by an evil demon). Then the result is this: $K\sim e.d. \ \& \ K(e.d.) \rightarrow \sim e.d. \rightarrow K\sim e.d.$

¹⁴ Stroud's principle may have been motivated by the thought that in order to know that p , one must be able to *rule out* all alternatives that one knows to be incompatible with knowing that p – where *ruling out* means being able to

determine that those alternatives do not obtain *only on the basis of evidence that one would have even if they did obtain*. However, it is not true that in order to know that a possibility doesn't obtain one has to be able to rule it out in this sense, as is shown by the example regarding the children of brunettes discussed on p. [] below.

¹⁵ This thought will be unavailable if one understands of every episode of critical reflection or deliberation about a possibility on the model of performing an explicit inference. I'm trying to bring out that that is a mistake.

(A note for future thought. What am I doing when I provide a ground for dismissing a suggestion, if I am not thereby performing an inference or providing a reason or evidence on which I base my belief? Perhaps I am *explaining why* the suggestion is ridiculous by highlighting something which is the case and which decisively tells against its truth.)

¹⁶ This point will not be available if one simply assumes that there *must* be an epistemic priority relation here – i.e., that either the belief that one is not being deceived by an evil demon is based upon the belief that there are no evil demons, or the belief that there are no evil demons is (partially) based upon the belief that one is not being deceived by an evil demon; or at the very least, one of these propositions must come before the other in the “order of justification.” Part of what I am trying to bring out is that nothing in our ordinary commitments forces us to accept this requirement.

¹⁷ For the purposes of this argument, it doesn't matter what substantive theory of epistemic priority one chooses to work with. It suffices to treat the notion of “antecedently knowing” as an uninterpreted notion with a certain formal structure. (That is likewise all that would be needed in order to get the skeptical argument to work in the suggested way.)

¹⁸ In this paragraph and what follows, I mean the term “ground” to allow anything that one might want to allow as an epistemically acceptable basis for a belief. (I do assume, however, that practical or prudential reasons won't count.)

¹⁹ I have addressed Stroud's version of it in detail elsewhere [Leite 2005]. The seeds of that response are contained in what I say here in response to Bonjour.

²⁰ Suppose that R is not itself a consideration about the world around us and *is* a reason to believe something about the world around us. Could one answer Bonjour's question by first basing a belief about the world on R and then using that belief as a basis for concluding things about what reasons one has for beliefs about the world (including that R is such a reason)? Such a procedure would be cheating if it is intended as a response to the project Bonjour intended to initiate. One is supposed to determine what reason (if any) one has for beliefs about the world *in a way that does not depend upon any considerations about the world*. One is first supposed to bracket one's commitments about the world. One is then supposed to determine what reasons (if any) one has for believing things about the world. And one is finally supposed to regain those commitments about the world, now grounded upon what have been independently shown to be good reasons. So the proposed procedure would take things in the wrong order, given the aims of

the project. (And if the proposed procedure were acceptable, it would hardly help lay the ground for a *skeptical* argument.)

²¹ In a sense, we've come full circle. We began with the thought that mere assertion or consideration of the skeptical conclusion sufficed to do the trick, but it won't. Neither will mere consideration of the epistemologist's special question. Thanks to Katy Abramson for noticing this connection.

²² This point holds only so long as the "Moorean dogmatist" doesn't take the doctrine of immediate experiential justification to imply that a fully rational and reflective agent could acceptably reason from her sensory experiences to such conclusions as that she is not being deceived by an evil demon (see p. []). For an extended argument that the doctrine of immediate experiential justification does not require the further claim, see my [forthcoming2].