

TAKING SKEPTICISM SERIOUSLY
Adam Leite (Indiana University, Bloomington)
May, 2009

Modern-day heirs of the Cartesian revolution have been fascinated by the thought that one could 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 believe most if not all of what one currently believes about the world. A good part of the interest and mystique of these discussions arises from the contention that the seeds of such arguments lie in our ordinary epistemic practices, so that external world skepticism can arise “from within”. But is this contention correct? I doubt it.

Taking skepticism seriously requires that we address this question head on. To do so, 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 or 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. From within that vantage point can I somehow be moved in a reasonable way to accept the conclusion that I know 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 such path from our ordinary position to the skeptical conclusion. Here I will focus on perhaps the most prominent strategy of this kind. My goal is to show the fruitfulness of my approach by enacting it. The upshot is, in effect, a challenge. If you think you can get a global skeptical argument out of the commitments of ordinary life, show me the money. Otherwise, let's stop taking seriously the idea of global skepticism from within.¹

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".²

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.³

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. The same reasoning applies to each of them. So I should not believe any of those things.

These are relatively crude presentations, and there are lots of details about which one might fuss. But I'm presenting matters this way so that I can easily highlight a fundamental issue: 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?

For a parallel, consider the following argument:

- i) If we are in the twenty-first century, then the year is after 1999.
- ii) But the year is not after 1999.
- iii) So, we are not in the twenty-first century.

The correct response to this argument, given the principles and procedures of ordinary life, is this: “Valid argument, but with a clearly false premise.” It would be *crazy* to treat this argument as giving you some reason to accept its conclusion. My question, then, is simply this: from the standpoint of ordinary life, why should I regard the skeptical arguments any differently?

Consider the position we were starting from when we encountered the skeptical 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.⁴ 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 standing in 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, standing in my ordinary position it seems that 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 know that according to the lore of our tradition, I’m not supposed to respond in this way. But I want to know why not. Why can’t I make use of considerations about the world at this juncture? All the crucial work must actually be done here: in explaining what’s wrong with this sort of response. 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. Following Stroud, 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.⁵ I will focus on issues having to do with the adequacy of our evidence, the suggestion that the skeptical hypothesis is constructed so that none of our ordinary evidence counts against it, 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 Some 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. The basic point I want to urge is one whose time has come.⁷

Begging the question, first pass

Let's begin with a simplistic version of the charge that my 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 this way, 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.

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 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 .⁸ But as the burgeoning literature on this topic has shown, the Sensitivity requirement is vulnerable to straightforward counterexamples. Here's a simple one, inspired by Timothy Williamson.⁹ 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.¹⁰

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, as the possibility of knowledge on fallible evidence shows. Suppose 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. I would not possess that evidence if I were being deceived by an evil demon, since in that case it would not be true that there are no evil demons. 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.¹¹

Of course, if I were being deceived by an evil demon I would incorrectly believe that there are no evil demons. 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: when 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. 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?

Following the lines of a proposal of Ram Neta's, 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*. When one considers such a hypothesis, according to this proposal, one thereby restricts what counts as *S*'s body of evidence to just the evidence *S* would have whether or not *H* is true.¹² This 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, however. 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.

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 a key burden of Stroud's argument in chapter 1 of *The Significance of Philosophical Scepticism*.¹³) 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 global skeptical arguments. 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 one's beliefs about the world.¹⁴ 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.¹⁵ 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.¹⁶

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 .¹⁷ (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.¹⁸

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.

- a. 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 of A.I.)
- b. 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 (a) and (b) 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 *requires* our knowing other things about the world.

An example may help here. Consider the following two things I know: 1) I am on the graduate faculty of my university; 2) I have a standing entitlement to serve on PhD dissertation committees at my university. It's perfectly possible for someone to come to know an analog of one of these (regarding him/herself or someone else) without knowing the other. But given what I know about the relation between these (you can't have the one status without the other), my knowing either requires knowing the other. There is obviously no problematic circle in this situation. There is no reason to think that either of these beliefs is the ground on which I base the other. (For instance, order of acquisition is irrelevant to this question, and it may be that I found out about the relation between these two only after I came to know each of them. And standard counterfactual tests won't establish that one of these items of knowledge is the ground on which I base the other: if I stopped believing one, I'd stop believing the other too.) Moreover, there's no reason to think that there is some other sort of priority relation between these belief-contents, such that my beliefs will not be in good order unless one of these two beliefs is grounded on the other but not vice versa. So this example shows us that there can be mutual epistemic dependence without objectionable epistemic circularity.

A skeptical conclusion would result if (a) or (b) above 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 (a) or (b) 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 know that I am not being deceived by an evil demon. And that *is* an unsatisfiable requirement.

Now it might seem that I have already committed myself to a priority reading of (b). 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 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 that there are no evil demons I would already have to know 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. 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: there's obviously nothing problematic about my offering *both* of these responses. So one can appropriately provide a ground for dismissing a possible hypothesis without thereby creating or tracking an epistemic priority relation.¹⁹

It might be objected that I'm still not out of the woods, however, because my belief to the effect that I am not deceived by an evil demon is based on other beliefs about the external world, and so vicious circularity is bound to arise from my accepting both (a) and (b). However, I can concede to this objector all of the

following: (1) Beliefs about the world played a causal role in my coming to believe that I am not being deceived by an evil demon, even perhaps a causal role that a cognitive scientist might want to model for theoretical purposes as an implicit inference, (2) I believed things about the world before I believed that I was not being deceived by an evil demon, (3) I take some of the things I believe about the world to tell decisively in favor of the truth of the claim that I am not being deceived by an evil demon, (4) if I were asked to state considerations that provide good reasons for believing that I am not being deceived by an evil demon, I would appeal to considerations about the world, (5) if I stopped believing some of the things I believe about the world, then (being rational) I would stop believing that I am not being deceived by an evil demon (because I would become agnostic on the issue). Conceding all of that does not lead to vicious circularity if we add that my now knowing those other things about the world requires me to know that I am not being deceived by an evil demon.

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 other beliefs about the world, or vice versa, or at the very least one or the other of these must come before the other in the “order of justification.” Nothing in our ordinary commitments forces us to accept this requirement, however. So far, then, no reason has emerged for thinking that the conjunction of (a) and (b) is problematic.

I've been focusing on (b). Let's look at (a). Suppose that (a) – that is, premise (1) of the Argument from Ignorance – involved an epistemic priority condition. From that premise there would then 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 wreak havoc with our ability to arrive at knowledge through deductive inference. Even more radically,

however, it would straightforwardly make all knowledge impossible for anyone with minimal logical acumen, as becomes evident as soon as you substitute not-q for p into the schema and then consider the upshot for your knowledge of any p that you recognizes to imply not-q.²⁰ The same points apply to Stroud's variant of the closure principle as well. So this is no way to show that the priority version of premise (1) is a commitment of our ordinary practice.²¹

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

Here's a more promising strategy for deriving something like a priority reading of premise (1), a strategy that picks up on the common idea that the skeptical hypothesis is constructed in such a way as to "neutralize" our ordinary evidence. 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 from outside that range for believing that I am not being deceived by an evil demon. And if that's so, then in appealing to considerations about the world in response to the evil demon hypothesis, I wouldn't be showing that I have adequate grounds from outside that range. Maybe that's what's wrong with my response to the skeptical argument.²²

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 domain 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 a certain kind of epistemic priority structure between the belief that not-H and considerations in the domain D: one cannot acceptably believe anything in the domain without having adequate grounds from outside the domain for believing not-H. (I will use the term “excluded domain” to refer to this domain.)

Applying this schema to the evil demon hypothesis, 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.

Could this requirement explain why we can’t reject the evil demon hypothesis simply by appealing to 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, for the first time and 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 excluded 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 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 Bush family 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, imagine someone who does not yet have any relevant information or 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 Bush. They observe that Bush looks like a man, etc. 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. The same goes for someone who consciously and explicitly engages in a one-step inference from (1) to (3) in order to determine whether Bush is a woman engaged in a massive deception.²³

I consequently conclude that in order to believe such things as that Bush is a man, 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, for we have no such evidence. This version of the requirement-schema should consequently be rejected.

Let's pursue the argument one step further. It might be urged that Bush's looking like a man constitutes independent grounds that *do* decisively tell against 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 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 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 for the first time at a belief that the hypothesis is incorrect through an explicit, conscious course of 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 brunet 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 brunet 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. I've urged that this version of the requirement is incorrect and so won't do the work needed by the skeptic. But even if I'm wrong that this version of the 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-brunets hypothesis in the relevant

respect. I do not think this can be done. It might be suggested that the distinction can be drawn 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-brunets hypothesis as well. *Of course* I can acceptably reject out of hand the hypothesis that children of brunets 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 to a position in which we deem it inappropriate to rely upon

considerations about the world when determining what we know or ought to believe.

“Internalist” Constraints?

It is sometimes suggested that certain general features of our ordinary epistemic practices create an inhospitable environment for a response to the skeptical argument that merely appeals to considerations about the world. In particular, it is sometimes suggested that we especially value “reflectively accessible evidence” – evidence that one can possess and identify as such when one initially brackets all of one’s convictions about the world and utilizes only the resources of introspection and a priori reflection. And it is often thought that the demand for such evidence underwrites “internalist” conceptions of knowledge, justification, or other terms of epistemic approbation that would block any response to the skeptical argument that merely appeals to our ordinary conviction that there are no evil demons or to any other claim about the world around us: we won’t have fully dealt with the evil demon hypothesis unless we have provided “reflectively accessible evidence” against it. Of course, I haven’t done that. The question is whether anything in our ordinary epistemic practices requires me to.

In thinking about this issue, it's important to keep in mind that epistemological usage of the word "justification" is theory driven. Epistemologists use the term to capture a variety of statuses and potential features of beliefs, belief-contents, or believers that are of significance for their particular theoretical purposes, and so consulting "intuitions" about what is or isn't a case of justified

belief is pretty much worthless. We have to scrutinize what is meant in particular cases, and we have to ask how the status or features so picked out relate to the various statuses and distinctions that are of significance in our ordinary practice.

Consider, for instance, a classic argument of Chisholm's. Chisholm's premise is this:

(P2) I am justified in believing that I can improve and correct my system of beliefs. I can eliminate the ones that are unjustified and add others that are justified, and I can replace less justified beliefs [by more justified beliefs].²⁴

From this we are to conclude that a particular conception of justification is in play such that one "can know directly, just by reflecting on the state of his mind, that he is justified in believing that p" – where this is meant to exclude knowing in a way that involves relying on considerations about the world around us.²⁵

We can grant that Chisholm's premise is true: a key orienting principle of our epistemic practices is that it is possible for people to identify beliefs that are epistemically badly off and replace them with beliefs that are better off. But this does not commit us to thinking that we can know *directly, just by reflecting on the state of our minds*, whether a belief is badly off or not. It's perfectly compatible with this aspect of our practice to hold that we identify and eliminate particular beliefs that are epistemically badly off by making use of other things we know about the world. So this aspect of our practices does not generate a requirement that limits our evidence in a way that would warrant dissatisfaction with my reply to the skeptical argument.

Similar problems plague every familiar argument that our ordinary practice commits us to such a requirement. One commonplace in defenses of “internalism” is to appeal to considerations about epistemic responsibility, where responsible belief is taken to require that one have (and recognize) some reason for thinking that one’s belief is true.²⁶ But epistemic responsibility, so understood, would not impose the limitation at issue here; for all it demands, one’s reasons could perfectly well involve other claims about the world in every case. Something additional must therefore be brought into play to impose the proposed limitation. A standard argument appeals to the thought that it is constitutive of justification (or of evidence or reasons) that considerations about it can play a regulative role in deliberations about what to believe. But we can perfectly well regulate our beliefs in the light of something that is not accessible purely through introspection and a priori reflection.²⁷ One would only think otherwise if one were imagining that the relevant sort of regulation must begin from a position in which one had no empirical beliefs about the world or in which all such beliefs were “off the table.” But that conception does not fall out of the idea that justification, evidence, or reasons are “regulative”. Rather, it is of a piece with the limitation that the argument was supposed to ground.

Another standard argument attempts to derive “internalist” requirements from the thought that our ordinary epistemic practices involve considerations about epistemic duties or obligations. Since one cannot be held responsible for meeting or failing to meet one’s duties unless one is able to find out what they are, the factors determining one’s epistemic duties in a particular case must be

something one can find out about. However, it is compatible with this requirement that in certain cases one's ability to find out one's duty might ineliminably depend upon considerations about matters that are not accessible through introspection and a priori reflection alone. One will think otherwise if one thinks that it must be possible to find out about one's epistemic duties even if one had no empirical beliefs about the world or if all such beliefs were "off the table." But, again, that conception is of a piece with what the argument was supposed to establish.

It has been suggested that epistemic rationality or justification is a kind of instrumental rationality: it is epistemically rational for one to believe a given proposition just if one would, after careful reflection, take believing that proposition to be an effective means to one's overall epistemic goal.²⁸ However, this view would not limit one's evidence or reasons – or the factors relevant to rationality – to what one can find out about through introspection and a priori reflection alone. For it is perfectly compatible with this view that one's means-ends reflections would involve considerations about the world that one could not find out about in this way.

The problem for my objector here isn't just a lack of internalism-generating principles. There aren't any examples that appropriately connect with our ordinary practice, either. Some epistemologists have argued that examples of "chicken sexers" and the like – people able 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.”²⁹ But these examples don’t show any such thing. If we ordinarily judge such people 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.

Another example that has loomed large is that of the “new evil demon” argument.³⁰ Consider someone who is your epistemic twin being deceived by an evil demon: although this person is a disembodied spirit being deceived by an evil demon, everything experientially, introspectively, and intellectually seems the same to this person, and this person forms all of the beliefs that you do. Isn’t there some important epistemic status that the two of you share? The “internalist” says, “Yes,” and I see no reason to disagree. But even if we grant that this shared status shows that you and your twin share some evidence or reasons of a special “internal” kind, this concession doesn’t force any “internalist” limitation. The fact that you possess the status (and any evidence) that your twin possesses doesn’t entail that your twin possesses every status (and all the evidence) that you possess. Still, however, I see no reason to make the concession. The example itself doesn’t give us reason to think that a special “reflectively accessible” kind of evidence or reasons are in play. The shared status can be accounted for entirely in terms of the fact that both you and your twin blamelessly and responsibly take certain things to be evidence or good reasons for beliefs about the world. It might be said that your twin’s beliefs about the world “fit with” its seeming experiences of the world. But that thought can be explained without supposing that the twins’

experiential seemings are good evidence or reasons that can be identified as such using only the resources of introspection and a priori reflection.

It is sometimes suggested that what are in effect internalist limitations are forced on us by certain very general questions about what we know, have reason to believe, or should believe.³¹ Consider, for instance, the question, “Do I have any good reasons for believing anything at all about the world around me?” Let’s concede that this is an intelligible question and that we are precluded from making ineliminable use of considerations about the world in answering it (because doing so would assume an answer to the very question under investigation). These concessions do not commit us to an “internalist” constraint to the effect that epistemic reasons must be available and identifiable as such when one does not rely at the ground-level upon considerations about the world. To see why not, suppose that you sincerely attempt to answer the general question about reasons and find that given only the limited resources permitted to you, you cannot reach a positive conclusion. Would you be entitled to conclude that you *do not* have any reason to believe anything about the world? No. It would remain possible that you have good reasons for believing particular things about the world and even that you can find out what they are. Perhaps you just have to proceed piecemeal, relying in each particular case on other considerations about the world.³² If an “internalist” constraint on reasons were forced merely by our concessions, you should have been able to reach the negative conclusion straight off. Since you can’t, our concessions don’t force an “internalist” constraint. The

same point applies to the general questions about knowledge or about what, if anything, we should believe.

Demands for “reflectively accessible” evidence for beliefs about the world – or considerations about the possibility of such evidence – are consequently irrelevant to my response to the skeptical argument. Admittedly, many “internalists” have argued that we do have good “reflectively accessible” evidence for beliefs about the world, including the belief that we are not being deceived by an evil demon.³³ But from the standpoint of our ordinary practices my response to the skeptical argument does not depend upon their success.

Postlude

I want to close by briefly reflecting on the overall perspective enacted here.

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 response to skepticism in “Proof of an External World,” and they may have seemed similarly dogmatic.³⁴ However, I was looking at an earlier stage in the skeptical argument and asking how the crucial first steps of that argument are 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 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*: roughly, one's sensory experiences can give one justification to believe something in a way that does not constitutively depend upon one's having any other justified beliefs or justification for any other beliefs.³⁵ The arguments that I offered are compatible with this theory.³⁶ 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 *prima facie* 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 about our practice. 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.³⁷

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

² The label comes from Keith DeRose's "Solving the Skeptical Problem," *The Philosophical Review*, 104 (1995), 1–52.

³ This presentation will idealize away from possible concerns about degrees of belief or credence and work with an on/off notion of belief.

⁴ For a full-dress discussion of this form of response to the dream-argument for external world skepticism, see my "Dreams and Skepticism: What Austin Saw and Moore Missed," forthcoming in Sorli and Gustafsson, eds., *New Essays on the Philosophy of J. L. Austin*, Oxford University Press.

⁵ Stroud, *Significance of Philosophical Scepticism* (New York, NY: Oxford University Press, 1984), chapter 1.

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

⁷ Others who I take to be heading in a similar direction are Ernie Sosa, Alex Byrne, and Penelope Maddy. See especially Sosa's *A Virtue Epistemology* (New York, Oxford University Press, 2007) and *Reflective Knowledge* (New York, Oxford University Press, 2009), Byrne's "How Hard are the Sceptical Paradoxes?," *Noûs* 38 (2004), 299-325, and Maddy's *Second Philosophy*, (New York, Oxford University Press, 2007, especially chapter 1). In their own ways G. E. Moore, J. L. Austin and W.V. O. Quine beat us to the punchline long ago.

⁸ Robert Nozick, *Philosophical Explanations* (Harvard University Press, 1981).

⁹ *Knowledge and Its Limits* (New York: Oxford University Press, 2000).

¹⁰ In Nozick's formulation of the Sensitivity Requirement (p. 179), the way or method by which the belief was formed must be held fixed across the actual and counterfactual cases. It might be objected that in my example the method of belief formation isn't held fixed, because one's perceptual experiences would be quite different in the two cases. However, there are three points to be made here. First, there is a straightforward (not *ad hoc* or jury-rigged) characterization on which the method of belief formation is the same in the two cases. Second, Nozick's own examples and his explanation of why, e.g., one's current belief that one has hands meets the Sensitivity Requirement, allow for differences of this sort between the cases. Third, if we required that the relevant counterfactual case be one in which one's current sensory experience is introspectively indistinguishable from the actual case but one's belief is false, then *many* of our ordinary beliefs that we take to constitute knowledge wouldn't meet the Sensitivity Requirement and so wouldn't count as knowledge.

¹¹ If there is a general principle at play in this argument, it is something along the following lines: if p is false, then p isn't evidence for anything. A principle along such lines accords with our ordinary practice. However, Branden Fitelson has offered a putative counterexample. Suppose that a reliable source tells me that p, and I thereby (reasonably) come to believe that p. Recognizing that p entails some q, I then reasonably infer q from p and thereby take p to be evidence for q in this case. Branden comments, "I'm inclined to say that p could *be* evidence for q (for me) in some such contexts – even if it turns out that p is false (especially, if I'm – in principle – never in a position to know that p is false)" (personal communication; Jon Kvanvig has suggested something similar to me as well). However, there is a slide away from ordinary ways of thinking if one takes such examples in the suggested way. We should grant that one could reasonably take p to be evidence in this sort of case. But note the shift in terminology – the relativization to a person – involved in the suggestion that p could "be evidence for q (*for me*)" in such a case. This may well just mean the same thing as "S reasonably takes p to be evidence for q" or "S would be reasonable to treat p as evidence for q" -- which is compatible with p's *not* being evidence for q. But even if that is not what's meant, such formulations are not equivalent to saying that p is evidence for q simpliciter. And we wouldn't ordinarily say things like, "Well, p is evidence for S, but not for me, since (I know) p is false." Rather, we'd say, "S reasonably thinks p is evidence for q, but he's wrong." I conclude that *evidence-for-a-person* is a technical, theoretical notion aimed at helping us explain the sense in which someone can be reasonable in certain sorts of cases.

Confusion between this technical notion and our ordinary notion of evidence is facilitated, I think, by a commitment to the thought that if S's belief is reasonable, then it must be based on good evidence, or at least S must have good evidence for it. But the judgments involved in our ordinary practice do not themselves show a commitment to that thought.

¹² 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 own formulation is considerably more complicated, for reasons arising from the details of his theory of evidence.

¹³ New York: Clarendon Press of Oxford University Press, 1984.

¹⁴ Op cit., p. 25.

¹⁵ Op cit., pp. 19 ff.

¹⁶ Let $p = \sim ED$. (the evil demon hypothesis is not the case). Let $q = ED$. (I am being deceived by an evil demon). Then the result is this: $K\sim ED \ \& \ K(\sim ED \rightarrow \sim ED) \rightarrow K\sim ED$

¹⁷ Op cit., pp. 29-30.

¹⁸ 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.

¹⁹ This thought will be unavailable if one understands 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.

²⁰ 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.

The implausibility of this principle has also been noted by James Pryor ("The Skeptic and the Dogmatist," *Noûs* 34, 517-49) and Byrne, *op cit*. I first articulated the problem in my undergraduate senior honors thesis (UC Berkeley, 1992).

²¹ 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.)

²² In this paragraph and what follows, I mean the term “ground” to include 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.)

²³ There is now an extensive literature concerning what exactly goes wrong in instances of reasoning like this one. On my view, what goes wrong here has to do with the commitments that are undertaken when one forms a belief by consciously and explicitly inferring from some evidence (see my “Believing One’s Reasons are Good,” *Synthese*, 161(3), 2008, 419-441 and “Immediate Warrant, Epistemic Responsibility, and Moorean Dogmatism,” *Synthese* (forthcoming)). For some other approaches, see Crispin Wright, “Some reflections on the acquisition of warrant by inference” in S. Nuccetelli, ed., *New Essays on Semantic Externalism and Self-Knowledge*, (Cambridge, MA: MIT Press, 2003), 57-77, “Warrant for nothing (and foundations for free?)” *Aristotelian Society Supplement* 78, 2004, 167-212, and “The perils of dogmatism,” in *Themes from G. E. Moore: New Essays in Epistemology and Ethics*, ed. by Susana Nuccetelli and Gary Seay (New York, NY: Oxford University Press, 2007), and the literature regarding Stewart Cohen’s so-called “problem of easy knowledge” (Cohen, “Why Basic Knowledge Is Easy Knowledge” *Philosophy and Phenomenological Research*, 70(2), March 2005, 417-430 and “Basic Knowledge and the Problem of Easy Knowledge,” *Philosophy and Phenomenological Research*, 65(2) , September 2002, 309-329). (For a dissenting voice, see Jim Pryor, “When Warrant Transmits,” *Wittgenstein, Epistemology and Mind: Themes from the Philosophy of Crispin Wright*, ed. by A. Coliva, Oxford University Press (forthcoming) and “What's Wrong with

Moore's Argument" *Philosophical Issues* (a supplement to *Nous*), 14, 2004, 349-378.)

²⁴ Chisholm, R. M., *Theory of Knowledge*, 3rd ed. (Englewood Cliffs: Prentice-Hall, 1989), p. 5.

²⁵ *Ibid.*, p. 7.

²⁶ See, for instance, Laurence Bonjour, *The Structure of Epistemic Justification* (Cambridge, MA: Harvard University Press, 1985), chapter 1.

²⁷ Alvin Goldman, "Internalism Exposed," *Journal of Philosophy* 96:6, 1999, 271-293.

²⁸ R. Foley, *The Theory of Epistemic Rationality* (Cambridge, MA: Harvard University Press, 1987), chapter 1.

²⁹ E.g., Duncan Prichard, *Epistemic Luck* (New York, NY: Oxford University Press, 2005), 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.

³⁰ The example first appeared in Cohen, S. and K. Lehrer, "Justification, Truth, and Knowledge," *Synthese* 55 (1983) and Cohen, S., "Justification and Truth," *Philosophical Studies* 46 (1984), 279-96.

³¹ See, for instance, Stroud, "Understanding Human Knowledge in General," in *Understanding Human Knowledge: Philosophical Essays* (Oxford: Oxford University Press, 2002) and Bonjour, *Epistemic Justification: Internalism vs. Externalism, Foundations vs. Virtues* (with E. Sosa, Oxford: Blackwell, 2003), chapter 1. Stroud focuses on the question, "How does anyone know anything

about the world at all?”, whereas Bonjour focuses on the question, “Do we have any good reasons for thinking that our beliefs about the world ... are true?”

Another question that would have much the same effect is, “What, if anything, am I to believe about the world?”

³² For a parallel, suppose that someone asked you to find out, by sight, how many items of furniture are in a completely dark room. When you fail, it would be a mistake to conclude that there is no furniture in the room or even that you can’t find out whether there is. To find out whether there is, you’d either need to turn on the lights or use something other than your eyes. As with the analogy, so too regarding the question about reasons.

³³ See, for instance, Bonjour, *Structure of Empirical Knowledge*, op. cit., pp. 179 - 188 and *Epistemic Justification: Internalism vs. Externalism, Foundations vs. Virtues*, op. cit., pp. 92-6.

³⁴ “Proof of an External World,” in T. Baldwin (ed.), *G. E. Moore: Selected Writings* (London: Routledge, 1993). See also Moore’s “A Defence of Common Sense,” same volume.

³⁵ Pryor, op cit.

³⁶ Compatible, that is, so long as one doesn’t follow the “Moorean Dogmatist” in assuming that there would be nothing epistemically objectionable about a fully rational and reflective deliberator consciously and explicitly inferring from her sensory experiences to such conclusions as that she has hands and from there that she is not a disembodied spirit deceived by an evil demon – thus arriving at a belief on the matter for the first time. (I argued above that parallel reasoning,

exhibiting the same structure, would be unacceptable in the George W. Bush case (see p. [] above.) However, the doctrine of immediate experiential justification does not force the additional claim about acceptable inference or reasoning. See my “Immediate Warrant, Epistemic Responsibility, and Moorean Dogmatism” (forthcoming).

³⁷ I would like to thank Gary Ebbs, Branden Fitelson, Mark Kaplan, Penelope Maddy, Berislav Marusic, Ram Neta, Fred Schmitt, and Jonathan Weinberg for comments and conversations relating to this paper. I am grateful, too, to the audience at an invited paper session at the 2009 Pacific Division American Philosophical Association meeting and to my commentators, Al Casullo and Jon Kvanvig, and session chair, Peter Graham. My greatest debt is to Katy Abramson, without whom these thoughts – some fifteen years in the gestating – would never have made it onto paper.