

The Skeptic and the Dogmatist

JAMES PRYOR
Harvard University

I

Consider the **skeptic about the external world**. Let's straightaway concede to such a skeptic that perception gives us no conclusive or certain knowledge about our surroundings. Our perceptual justification for beliefs about our surroundings is always defeasible—there are always possible improvements in our epistemic state which would no longer support those beliefs. Let's also concede to the skeptic that it's metaphysically possible for us to have all the experiences we're now having while all those experiences are false. Some philosophers dispute this, but I do not. The skeptic I want to consider goes beyond these familiar points to the much more radical conclusion that our perceptual experiences can't give us any knowledge or even justification for believing that our surroundings are one way rather than another.

One might go about grappling with such a skeptic in two different ways.

The **ambitious anti-skeptical project** is to refute the skeptic on his own terms, that is, to establish that we can justifiably believe and know such things as that there is a hand, using only premises that the skeptic allows us to use. The prospects for this ambitious anti-skeptical project seem somewhat dim.

The **modest anti-skeptical project** is to establish to *our* satisfaction that we can justifiably believe and know such things as that there is a hand, without contradicting obvious facts about perception. This is not easy to do, for the skeptic can present us with arguments from premises we find intuitively acceptable to the conclusion that we *cannot* justifiably believe or know such things. So we have a problem: premises we find plausible seem to support a conclusion we find unacceptable. The modest anti-skeptical project attempts to diagnose and defuse those skeptical arguments; to show how to retain as many of our pre-theoretical beliefs about perception as possible, without accepting the premises the skeptic needs for his argument. Since this modest anti-skeptical project just

aims to set our own minds at ease, it's not a condition for succeeding at it that we restrict ourselves to only making assumptions that the skeptic would accept.

Our next character is the **fallibilist**. A fallibilist is someone who believes that we can have knowledge on the basis of defeasible justification, justification that does not *guarantee* that our beliefs are correct. We can at best have defeasible justification for believing what our senses tell us; so anyone who thinks we have perceptual knowledge about our environment has to embrace fallibilism. I assume that most of us *are* fallibilists. Most of us think we *do* have some perceptual knowledge. Different fallibilists will reply to the skeptic's arguments that we can't know or justifiably believe anything on the basis of perception in different ways.

The next character in our story is **G.E. Moore**. Moore is famous for two anti-skeptical ideas.¹

Moore's first idea is that the proposition that there is a hand, though only defeasibly justified, is more certain than any of the premises which might be used in a skeptical argument that he does not know that there is a hand. This first anti-skeptical claim of Moore's sounds like a mere expression of confidence that the modest anti-skeptical project I described can be completed, somehow or other. Moore hasn't yet given us any instructions as to *how* to defuse the skeptic's argument. He's just observed, rightly enough, that if *something* has to give, it will be much more reasonable to reject the skeptic's premises than it will be to accept the skeptic's conclusion.

Moore's second anti-skeptical idea is that he can know some propositions without being able to "prove" them. Recall Moore's notorious proof of the external world: "Here is one hand, and here is another; hence there are external objects." Moore says that this proof is perfectly satisfactory because *he knows its premises to be true*, even though, as he admits, he is not able to "prove" them.² What's involved in knowing a premise to be true without being able to prove it? What sort of "proof" is Moore unable to give? Of course, Moore can't offer a *deductive proof* that he has a hand, from premises of whose truth he is more certain, and which the skeptic will accept. But it seems that his position is even worse than that. For Moore doesn't seem ready or able to offer *any considerations at all* in favor of the claim that he has a hand—even defeasible, ampliative considerations—without begging the question against a skeptic who refuses at this stage of the dialectic to grant the existence of the external world. This is why Moore's "proof" strikes us as so unsatisfactory: he hasn't offered any non-question-begging reasons to believe his premises. Yet Moore claims he can know these premises to be true. He can know them to be true, though he has no non-question-begging arguments to offer in their support.³

In this essay, I will formulate and defend an epistemology of perception based on this second anti-skeptical idea of Moore's. I call this position **dogmatism**.

Let me begin with a few words about the view of experience I'm working with. Epistemologists sometimes think of experience as entirely subjective or sensational in character. By contrast, contemporary philosophers of mind mostly

think of perceptual experiences as states with propositional content. Your experiences represent the world to you as being a certain way, and the way they represent the world as being is their propositional content. This propositional content is present to your mind simply by virtue of your having the experience, independently of any beliefs you might have about what external states of affairs the experience is reliably connected with. That is the model of experience I will be employing.⁴

Now I can say what a dogmatist epistemology amounts to. One can be a dogmatist about either perceptual justification or perceptual knowledge. The dogmatist about perceptual justification says that when it perceptually seems to you as if p is the case, you have a kind of justification for believing p that does not presuppose or rest on your justification for anything else, which could be cited in an argument (even an ampliative argument) for p . To have this justification for believing p , you need only have an experience that represents p as being the case. No further awareness or reflection or background beliefs are required. Of course, other beliefs you have might defeat or undermine this justification. But no other beliefs are required for it to be in place.

Note that the dogmatist is not saying that your justification for believing p rests on your *awareness* of your experiences. His view is that you have justification for believing p simply in virtue of *having* an experience as of p . On his view, your experiences give you justification for believing p , but it would be misleading to call these experiences your “evidence” for believing p . For saying that your experiences are your “evidence” for a perceptual belief suggests that your justification for that perceptual belief depends in part on *premises* about your experience—as if you were introspectively aware of your experiences, and your perceptual belief were based in some way on that awareness. The dogmatist denies that you need any “evidence” of that sort for your perceptual beliefs. Compare: when you have a justified belief that $p \ \& \ q$, you are thereby also justified in believing p . But this justification for believing p does not rest on any *awareness* you may have of the fact that you have a justified belief that $p \ \& \ q$. You do not need to be able to appeal to the fact that you have a justified belief that $p \ \& \ q$ as a *premise*. The *mere having* of a justified belief that $p \ \& \ q$ is enough for your justification for believing p to be in place. Similarly, the dogmatist thinks that the *mere having* of an experience as of p is enough for your perceptual justification for believing p to be in place. You do not, in addition, have to be *aware* of your experiences and appeal to facts about them as “evidence” for your perceptual beliefs.⁵

Of course, you can *become* aware of your experiences, by introspection. And your introspective awareness that you have experiences of certain sorts might, together with appropriate background beliefs, provide you with *additional* reason to believe p . The dogmatist does not deny that. He allows that you may have some justification for believing p that does rest on your introspective awareness of your experiences, and on background beliefs. He only claims that there is a kind of justification you have which *does not* rest on these things. This jus-

tification would be in place even if you lacked the other reasons—even if you were in no position to provide any non-question-begging argument for your perceptual belief.

That was dogmatism about perceptual justification. The dogmatist about perceptual *knowledge* adds the further claim that this justification you get merely by having an experience as of *p* can sometimes suffice to give you knowledge that *p* is the case.

It's important to understand the difference between dogmatism and other fallibilist epistemologies. The dogmatist about perceptual knowledge is a fallibilist, but few fallibilists are dogmatists.

Most fallibilists concede that the ambitious anti-skeptical project is hopeless: we can't demonstrate to the skeptic, using only premises he'll accept, that we have any perceptual knowledge. But that concession does not yet make one a dogmatist. The ordinary fallibilist thinks that having perceptual knowledge about the external world *does* require one to be in a position to provide some non-question-begging considerations in support of one's perceptual belief. When he concedes the ambitious anti-skeptical project, such a fallibilist is only acknowledging that the considerations which support our perceptual beliefs are defeasible and ampliative. The skeptic can grant these considerations (since they're not question-begging); but they don't *entail* that there's a hand, and the skeptic furthermore doesn't see why they should offer the conclusion that there's a hand any support whatsoever. There may be no way to *rationaly force* the skeptic to grant our perceptual beliefs, even if he accepts the evidence we cite in their favor. On the ordinary fallibilist's view, that is why the ambitious anti-skeptical project can not succeed.

The dogmatist's view is more radical than the ordinary fallibilist's. The dogmatist thinks that not only can we have perceptual knowledge and justified perceptual belief, we might have it without being in a position to cite *anything* that could count as ampliative, non-question-begging evidence for those beliefs.⁶ More on this as we proceed.

If we're going to give a proper reply to the skeptic, then we ought to begin by formulating the skeptic's reasoning in the most charitable way. I do not think that standard formulations of the skeptic's reasoning do this. In §II, I will explain why these standard formulations are defective, and I will propose new formulations of the skeptic's reasoning. I think these new arguments more accurately reconstruct the skeptic's reasoning, and they pose more formidable threats to our possession of perceptual knowledge and justified belief.⁷

Though they pose more formidable threats, I think these threats can be met. The second half of this paper is devoted to articulating and defending a dogmatist account of perceptual justification which enables us to reject the skeptic's arguments that we have no justified perceptual beliefs. I will present the positive case for my account in §III, and I will defend it against an important objection in §IV.

My primary concern will be questions about perceptual *justification* rather than questions about perceptual knowledge. This is because the connections between justification and knowledge are complicated. I believe that the account of perceptual justification I will be arguing for can be extended to provide an account of perceptual knowledge, as well. But I will not attempt to do that.

I will also be concentrating throughout on questions about *what propositions* we have justification for believing. When we're evaluating the epistemic merit of a belief you hold, we're sometimes concerned with more than just whether it's a belief in some proposition you have justification for. We also take into consideration what sorts of justification, if any, your belief is *based on*. One may have very good reasons for believing *p*, but base one's belief in *p* on bad reasons. In such cases, we sometimes say that *you* are unjustified in believing *p*, or that *your belief* is unjustified or "ill-founded," though it's a belief in a *proposition* you have good reason to believe.

The question "What makes something a proposition you have good reason to believe?" and the question "What makes a belief in such a proposition well-founded?" are both legitimate and interesting epistemological questions. But I take the first to be more basic, and in this essay, I will only be concerned with that question. So when I ask questions like "What makes your perceptual beliefs justified?" I mean: what makes them beliefs in propositions you have justification for believing? I am not here concerned with the further question, what it takes for those beliefs to count as "well-founded."⁸

If we want to discuss skepticism about perceptual justification, there is an initial obstacle we have to confront: most contemporary presentations of the skeptic's reasoning are posed as challenges to the possibility of knowledge, not as challenges to the possibility of justified belief. Many philosophers believe that there *are* compelling skeptical doubts about the possibility of perceptual justification, too.⁹ In §II, I will argue that they are right. But it will take some work to see this. It is not immediately obvious what a compelling skeptical argument against the possibility of perceptual justification looks like. We will have to spend some time digging, examining the more familiar arguments about knowledge, before we'll be in a position to construct a compelling skeptical argument which deals with justification. This examination of skeptical arguments about knowledge is only heuristic, and will be confined to §II. Our primary target is the skeptic about perceptual justification.

So the plan is: we first examine skeptical arguments about knowledge. That will provide us with a bridge to a compelling skeptical argument about justification. Then I will present my positive account of perceptual justification, and show how it enables us to diagnose and defuse the skeptic's argument.

II

Nowadays, it's standard to present the skeptic's reasoning in something like the following form. Let "you are being deceived by an evil demon" mean that your

perceptual experiences are false appearances presented to you by an evil demon. (You need not *accept* the false appearances, to be “deceived” in this sense.) Let “you are in a position to know *p*” mean that you possess some justification *G* for believing *p*, and if you were to believe *p* on the basis of *G*, the belief so formed would count as knowledge. Then the skeptic’s argument goes:

- (1) You are not in a position to know you’re not being deceived by an evil demon right now.
- (2) If you’re to know anything about the external world on the basis of your current perceptual experiences, then you have to be in a position to know that you’re not being deceived by an evil demon right now.
- (3) So, by *modus tollens*, you can’t know anything about the external world on the basis of your current perceptual experiences.

Very roughly, premise (1) is motivated by the thought that no amount of perceptual experience could enable you to determine whether or not you’re being deceived by an evil demon, since you’d be having exactly the same experiences even if you were being so deceived. Premise (2) is usually motivated by appeal to some sort of Closure Principle. The skeptic claims that, if you’re to know that things are the way they perceptually appear, then—since things *wouldn’t* be that way if your experiences were false appearances presented to you by an evil demon—you must be in a position to know that your experiences are not false appearances presented to you by an evil demon.¹⁰

I do not believe that argument (1)–(3) is the most effective formulation of the skeptic’s reasoning. There are several reasons for this. One reason is that some philosophers refuse to allow the skeptic to use claims like “I can’t know I’m not being deceived” as *premises* in his reasoning. Maybe skeptical argument can *convince us* that we can’t know we’re not being deceived; but why should we grant such a claim as a *premise* in a skeptical argument?¹¹ Yet many of these philosophers will still acknowledge that the skeptic’s reasoning has some intuitive force. This suggests that, at the intuitive level, the skeptic’s reasoning does not essentially require claims like “I can’t know I’m not being deceived” as premises. Because the argument (1)–(3) does rest on such a premise, it makes the skeptic’s case out to be less compelling, and easier to resist, than it otherwise might be.

Another problem with the argument (1)–(3) is that it does not generalize in the same ways that the skeptic’s reasoning intuitively seems to generalize. There appears to be a core structure of skeptical reasoning which is the same regardless of whether we’re talking about straightforward alternatives to what we purport to know, like being deceived by a demon, or about skeptical hypotheses like dreaming, which present obstacles to our acquiring knowledge that are not *incompatible* with what we purport to know. Similarly, the core structure of the skeptic’s reasoning appears the same regardless of whether we’re considering skeptical challenges to our possession of *knowledge* or skeptical challenges to

our possession of *justified belief*. We'd like a formulation of the skeptic's reasoning which makes this common core explicit.

The argument (1)–(3) does not do this. In the first place, Closure Principles tell us only that we have to know to be false things which are incompatible with *the truth* of what we purport to know. The appeal to such Closure Principles can't explain skeptical arguments that employ hypotheses like dreaming.¹²

In the second place, it's difficult to see how to convert the argument (1)–(3) into a compelling argument against the possibility of perceptual justification. Suppose we simply replaced all references to knowledge with references to justification. This would be the analogue of premise (1):

- (4) You can't be justified in believing you're not being deceived by an evil demon right now.

Why on earth should we accept this premise? Even if we can't *know* for sure whether or not we're being deceived by an evil demon, isn't it at least *reasonable* to assume that we're not being so deceived, absent any evidence to the contrary? There are important disanalogies between justified belief and knowledge which come into play here, and make (4) much less plausible than (1). For instance, justified belief is not factive; and hidden features of a belief's etiology aren't as obviously destructive of justified belief as they are of knowledge. (Any Gettier case attests to this. Standard Gettier cases just *are* instances of justified belief which fail to count as knowledge due to some fact about their etiology.) So the mere fact that one's belief is false and caused by an evil demon isn't enough to show that the belief is unjustified. *Even if one is being deceived by a demon*, it's possible to be justified in believing that one isn't. So why should we accept premise (4)? Why can't *we* be justified in believing we're not being deceived by a demon? Much further argument is needed before we ought to accept (4).

Because the argument (1)–(3) is unsatisfactory in this way, I will try to construct a better argument on the skeptic's behalf. This argument will not rely on any bald claim like the claim that you can't know you're not being deceived. It will also generalize in ways that the argument (1)–(3) does not.

The fallibilist who says he knows he's not being deceived by an evil demon will typically base this conclusion, in part, on things he knows about the external world on the basis of perception. He will say something like this:

Of course I know I'm not being deceived by an evil demon! I'm having experiences of hands and tables and the like, and, as I can see, there are in fact hands and tables all around me. So I'm not being deceived.

The fallibilist who says this will grant the skeptic that the reasons he has to believe he's not being deceived by an evil demon *beg the question* whether

there's an external world of the sort he purports to perceive. However, he does not think that impugns his reasons, since he's convinced that there is an external world and that he's perceiving it.

So long as we are willing to concede the skeptic that the reasons we have to believe we're not being deceived all rest on perceptual knowledge of our environment, then the skeptic can replace the bald claim (1) with the following premise:

- (5) Either you don't know you're not being deceived by an evil demon; or, if you *do* know you're not being deceived, it's because that knowledge rests in part on things you know by perception.

I think that this is a very plausible premise. The hypothesis that all of our present experiences are the deceptions of an evil demon is not absurd. It seems to be a genuine metaphysical possibility. So we can't reject that hypothesis out of hand. If we do know that we're not being deceived by an evil demon, it's plausible that that knowledge would have to rest on things we know about our environment on the basis of perception.¹³

There are some philosophers who reject even (5). They say we can know the demon hypothesis to be false without resting that knowledge on any of our perceptual knowledge about the environment. Some of these philosophers say there's *a generic rational presumption* that our senses are working reliably. Such a presumption can be defeated, but whenever it's not defeated, it gives us reason to believe that we're not being deceived. It might even enable us *to know* that we're not being deceived, without having to appeal to any of our perceptual beliefs. Other philosophers say that it's *a priori* unlikely that orderly, coherent courses of experience like the ones we actually have are the products of dreams or evil demons.¹⁴ Other philosophers say we have some good non-perceptual reason for believing that God exists, and that he wouldn't endow us with grossly unreliable senses, or permit us to be constantly deceived. Perhaps there are yet further ways we could come to know that we're not being deceived, without appeal to any of our perceptual beliefs.

I don't myself find any of these lines of argument fully convincing. So I'm going to grant the skeptic (5) for the sake of argument. My ultimate aim in this paper is to show that *even if we allow the skeptic that premise*, it is still possible to resist his central line of argument.

How would a skeptical argument which appealed to (5) rather than to (1) proceed? What the skeptic needs next is some reason to require us to know that we're not being deceived, and to know this *antecedently to* knowing anything on the basis of perception. The skeptic thinks that, in order for us to know anything on the basis of perception, we *first* have to know we're not being deceived.

What does this mean? When I speak of knowing one thing "antecedently" to knowing another, I'm not talking about temporal priority. Rather, I'm talking about epistemic priority. We can explain this relation as follows:

Your justification for believing p_1 is antecedent to your justification for believing p_2 just in case your reasons for believing p_1 do not *presuppose* or *rest* on your reasons for believing p_2 . Your reasons for believing p_1 can not *beg the question* whether p_2 .

For example, I might have good reasons for believing that the butler committed the murder. If my reasons for believing that the butler committed the murder crucially rest on the assumption that the murderer was left-handed, then I obviously couldn't defend *that assumption*, that the murderer was left-handed, by appeal to the claims that the butler was left-handed, and that *he committed the murder*. To do so would be question-begging. In this scenario, for me to be justified in believing that the butler committed the murder, I need to have some *antecedent* reason to believe that the murderer was left-handed, that is, some reason to believe that the murderer was left-handed which does not beg the question whether the butler committed the murder.

We can extend this notion of epistemic priority to knowledge as follows: you count as *knowing* p_1 antecedently to knowing p_2 just in case you know p_1 and p_2 , and *the justification* on which you base your belief in p_1 is antecedent to the justification on which you base your belief in p_2 .

Some cautionary remarks are in order. First, I do not mean to suggest that these relations of epistemic priority will be permanent or universal. *My* reasons for believing p might beg the question whether q ; whereas you have different sorts of evidence and so you have reasons for believing p which are antecedent to your reasons for believing q .

Second, your justification for believing one proposition can rest on your justification for believing other propositions, or beg the question whether those other propositions are true, even though you have not *formed* or *entertained* a belief in those latter propositions. Consider the following example:

I'm driving my car. I look at my gas gauge and it appears to read "E." This gives me justification for believing that I'm out of gas. However, for no good reason, I suspect that I'm hallucinating the gas gauge. So I do not actually form the belief that I'm out of gas, nor do I form the belief that my gas gauge reads "E."

In this example, since my suspicions are groundless and irrational, it seems plausible to say that I have justification for believing that I'm out of gas, even though I do not actually form this belief.¹⁵ What's more, there is an obvious sense in which my justification for believing that I'm out of gas *rests on* my justification for believing that the gas gauge reads "E"—even though I do not actually form either belief, and so do not *base* the one belief on the other.

Finally, admitting these relations of epistemic priority does not commit one to any strong form of foundationalism. Even moderate coherentists can allow that our reasons for believing *some things* are more basic than our reasons for

believing others. They deny that *the whole set of our beliefs* is organized in a foundational structure. That is compatible with allowing that *some* of our beliefs are more basic than others. (For instance, a coherentist can allow that, in the gas-gauge example, my reasons for believing that the gas gauge reads “E” are more basic than my reasons for believing that I’m out of gas. That does not mean that my reasons for believing that the gas gauge reads “E” are *absolutely* basic. The coherentist will argue that they in turn rest on reasons I have for believing that my senses are reliable. But it should be plausible, even for such a coherentist, that my reasons for believing that I’m out of gas rest in part on my reasons for believing that the gas gauge reads “E,” and that I’m not in a position to know the former unless I have antecedent reasons to believe the latter.)

I think that this notion of epistemic priority is a notion that we intuitively understand. It may be hard to *explain why* one’s opponent is begging the question, but it’s easy enough sometimes to see that he is and to understand the force of that criticism. And to say that one’s opponent is begging the question is just to say that he’s defending his conclusion with premises that *he lacks any antecedent justification for believing*. His grounds for believing those premises require him to *first* have reason to believe his conclusion.

I also think that this notion of epistemic priority is a notion that philosophers often implicitly appeal to. We’ll see later that it plays a crucial, but unacknowledged, role in Stroud’s discussion of the skeptic.¹⁶

I said that the skeptic needs some reason to require us to know that we’re not being deceived, and to know this antecedently to knowing anything on the basis of perception.

No Closure Principle of the standard sort will serve the skeptic’s purposes here. These principles are silent about knowing one thing antecedently to knowing another. Nor can the skeptic appeal to some Strengthened Closure Principle, like the following:

Strengthened Closure: If you know that *p* implies *q*, then you know *p* only if you are in a position to *antecedently* know *q*.

Such a Strengthened Closure Principle would be very implausible. Among other things, it would prevent us from ever acquiring knowledge by deductive inference.¹⁷

I think that, rather than appeal to some *general* reason for requiring us to know not-*q* antecedently to knowing *p*, whenever *q* is an alternative to *p*, the skeptic does better to appeal to the *special* features of his skeptical scenarios. The skeptic’s scenarios are not ordinary run-of-the-mill alternatives to what we purport to know on the basis of perception. They have special features. It’s these special features that account for the sense some philosophers have that no course of experience would enable us to know whether or not those scenarios obtain. The skeptic should argue that there’s something *especially bad* about the sce-

narios he puts forward, and that for this reason, we have to know that his scenarios do not obtain *antecedently* to knowing anything on the basis of perception.

Say that an alternative to p is a “bad” alternative just in case it has the special features that characterize the skeptic’s scenarios—whatever those features turn out to be. Different skeptical arguments will rely on different accounts of what makes an alternative “bad.” Here are some examples.

Say that some grounds E you have “allow” a possibility q iff the following counterfactual is true: if q obtained, you would still possess the same grounds E . Many skeptical scenarios are incompatible with what we purport to know on the basis of our experiences, but are “allowed” by those experiences, in this sense. For instance, your experiences at the zoo seem to justify you in believing that there is a zebra in the pen. This belief is incompatible with the hypothesis that the animal in the pen is a mule painted to look like a zebra. But that is a hypothesis which is “allowed” by your experiences: if it *were* a painted mule in the pen, you would most likely be having *the same experiences*, and hence, the same grounds for believing that there is a zebra in the pen. Likewise, the belief that there is a zebra in the pen is incompatible with the hypothesis that your experiences are false appearances presented to you by an evil demon; but this demon hypothesis is also “allowed” by your experiences. If it were to obtain, you’d be having exactly the same experiences. This is what tempts so many people to believe that they *can’t tell* whether or not the demon hypothesis obtains. So we might want to count a hypothesis as “bad” for the purpose of a skeptical argument just in case it is (and is recognized to be) incompatible with what you purport to know, but it is nonetheless “allowed” by your grounds E , in the sense I described.¹⁸

Some skeptical hypotheses don’t fit that paradigm. For instance, the skeptical hypothesis that you’re dreaming is *compatible* with many of your perceptual beliefs. You might still have hands even if you were dreaming right now that you have hands. The dreaming hypothesis does however introduce a non-standard explanation of your experiences. And this explanation would *undermine* the support your experiences give you for your perceptual beliefs—in the sense that, if you were to learn that you are dreaming, then you would have reason to doubt that your experiences were a trustworthy basis for beliefs about the external world. So we might want to count a hypothesis as “bad” for the purposes of a skeptical argument if it could undermine your experiences, in this sense.¹⁹ Both dreaming and being deceived by an evil demon are “bad” in this way.

For our discussion it does not matter which of these accounts of “badness” the skeptic adopts. Let’s suppose the skeptic does have some such account. His argument would then go as follows. From before, we have premise (5):

- (5) Either you don’t know you’re not being deceived by an evil demon; or, if you *do* know you’re not being deceived, it’s because that knowledge rests in part on things you know by perception.

We add two new premises:

- SPK If you're to know a proposition p on the basis of certain experiences or grounds E , then for every q which is "bad" relative to E and p , you have to be in a position to know q to be false in a non-question-begging way—i.e., you have to be in a position to know q to be false *antecedently* to knowing p on the basis of E .
- (6) The hypothesis that you're being deceived by an evil demon is "bad" relative to any course of experience E and perceptual belief p .

SPK stands for "Skeptical Principle about Knowledge." How the skeptic motivates this principle will depend on his choice of skeptical hypothesis and his account of what makes a hypothesis "bad." Suppose the skeptic does persuade us to accept SPK and (6). If we plug the demon hypothesis in for q ,²⁰ we get this:

- (7) If you're ever to know anything about the external world on the basis of your perceptual experiences, then you have to be in a position to antecedently know you're not then being deceived by an evil demon.

What follows from (5) and (7)? Let's suppose for *reductio* that you can know you're not being deceived. (5) says that knowing that you're not being deceived requires you to have some second piece of knowledge, which you got by perception, and on which your knowledge that you're not being deceived rests. Call this supposed second piece of knowledge p^* . Remember that p^* is supposed to be a piece of *perceptual* knowledge. (7) tells us that you can have perceptual knowledge of p^* only if you have some way of knowing you're not being deceived which *does not* rest on your knowledge of p^* . But, given our suppositions, you *don't* have any such p^* -independent way of knowing you're not being deceived. So it follows from (7) that you *cannot* have perceptual knowledge of p^* . Our supposition that you can know that you're not being deceived has now led to an absurdity: namely, that you both do and do not have perceptual knowledge of p^* . Hence, our supposition must be false. In other words:

- (8) You *can't* know you're not being deceived by an evil demon.

From (7) and (8), the skeptic can conclude:

- (9) You can't know anything about the external world on the basis of your perceptual experiences.

This then is the new skeptical argument I want to put forward. The argument is valid, and as I said before, premise (5) is very plausible. So everything will turn on the principle SPK.

I believe that this argument (5)–(9) reconstructs the skeptic’s reasoning more accurately than the argument (1)–(3). In addition, since this new argument does not rely on a bald claim like (1), but rather on the more plausible (5), it poses a compelling and formidable threat to our possession of perceptual knowledge. The skeptic still has work to do: he has to persuade us to accept SPK. We have not looked in any detail at how he might do that. But if you look at informal presentations of the skeptic’s reasoning, you’ll find that these *do* often rely on some principle like SPK, though that principle is hardly ever explicitly stated.

For instance, let’s look at part of Stroud’s discussion. Stroud rejects the **“straightforward response” to the skeptic** which (i) accepts that it’s a condition for having any perceptual knowledge that one know one is not dreaming, but (ii) maintains that one can meet this condition: one can know that one is not dreaming.²¹ Stroud is willing to entertain the possibility that we know we’re not dreaming, but he argues that we *wouldn’t* be able to know this if it were in fact a condition for having perceptual knowledge that one know one is not dreaming. That is, on Stroud’s view, parts (i) and (ii) of the “straightforward response” cannot both be correct. Stroud writes:

But how could a test or a circumstance or a state of affairs indicate to [Descartes] that he is not dreaming *if* a condition of knowing *anything* about the world is that he know he is not dreaming? It could not. He could never fulfill the condition... (p. 21)

In order to know that his test had been performed or that the state of affairs in question obtains Descartes would...have to establish that he is not merely dreaming that he performed the test successfully or that he established that the state of affairs obtains. How could that in turn be known? ...Some further test or state of affairs would be needed to indicate that the original test was actually performed and not merely dreamt, or that the state of affairs in question was actually ascertained to obtain and not just dreamt to obtain... And so on. At no point can he find a test for not dreaming which he can know has been successfully performed or a state of affairs correlated with not dreaming which he can know obtains. He can therefore never fulfill what Descartes says is a necessary condition of knowing something about the world around him. He can never know that he is not dreaming. (pp. 22–23)

Stroud’s argument in this passage seems to be:

- (10) Fulfilling the skeptic’s condition for perceptual knowledge (i.e., knowing that we’re not dreaming) would require us to have some piece of knowledge (e.g., that such-and-such a test has been performed successfully) which it would only be possible to have if the skeptic’s condition were fulfilled.

Hence, Stroud concludes that:

- (11) We can't fulfill the skeptic's condition.

Now, if the skeptic's condition is merely a *necessary condition* for perceptual knowledge, then Stroud's argument is not valid. With the skeptic's condition so interpreted, premise (10) is of the form:

- (12) For condition *C* to be fulfilled, some other fact *q* has to be known; and this fact *q* can be known only if *C* is fulfilled.

This does not entail that *C* cannot be fulfilled. What it entails is that *C* is fulfilled if and only if *q* is known.

For Stroud to have a valid argument that the skeptic's condition cannot be fulfilled, we have to understand the skeptic as imposing not a mere *necessary condition* for perceptual knowledge, but rather a *precondition* for perceptual knowledge.

With the skeptic's condition so understood, premise (10) should instead be read in this way:

- (13) For condition *C* to be fulfilled, some other fact *q* must *already* be known; and this fact *q* can be known only if *C* is *already* fulfilled.

If the "already" here signifies some asymmetric relation, then this does entail that *C* cannot be fulfilled. Hence, for Stroud's argument on behalf of the skeptic to succeed, we have to understand the skeptic as requiring us in some sense to *already* know we're not being deceived, if we're to have any perceptual knowledge. We can explain this in terms of the notion of antecedent knowledge I've sketched. Stroud's skeptic requires us to know we're not dreaming *antecedently* to having any perceptual knowledge. We have to know we're not dreaming in a way that doesn't beg the question whether any of the things we purport to know by perception are true. Such a skeptic is relying on SPK.²²

An advantage of the skeptical argument (5)–(9) is that it seems to translate into an equally compelling argument against the possibility of perceptual justification. If the argument works for knowledge, then it ought also to work for justification. Consider the following premise:

- (14) If you are justified in believing that you're not being deceived by an evil demon, that justification has to rest in part on some perceptual justification you have for believing things about the external world.

This premise just says that our reasons for believing we're not deceived are partly perceptual. Hence, it should be as plausible as its analogue (5) for knowledge was. Next, consider the analogue of SPK for justification:

SPJ If you're to have justification for believing p on the basis of certain experiences or grounds E , then for every q which is "bad" relative to E and p , you have to have antecedent justification for believing q to be false—justification which doesn't rest on or presuppose any E -based justification you may have for believing p .

For example, you have to have antecedent justification for believing that you're not being deceived by a demon, and that you're not dreaming, and so on. This does not mean that you have to *actually believe* you're not deceived (and so on), whenever you form a perceptual belief. Nor does it mean that you have to *infer* your perceptual belief from a prior belief that you're not deceived. So far as SPJ is concerned, your perceptual belief can be *spontaneous and uninferred*. It's just that, if that perceptual belief is to be justified, some antecedent reasons to believe you're not deceived have to be in place.

This principle SPJ is not obviously false. There is some plausibility to the idea that we're entitled to rely on our perceptual beliefs only if we're antecedently entitled to the assumptions that we're not deceived by an evil demon, or dreaming, and so on. Hence, although the details will depend on the skeptic's account of "badness," the major premises in the skeptic's argument (5)–(9) seem to be just as defensible when translated into a skeptical argument against the possibility of perceptual justification.

Another advantage of this form of skeptical argument is that, since it more accurately reconstructs the shared core of the skeptic's reasoning, it helps us to more clearly identify the places where that reasoning is vulnerable to attack.

We've already mentioned some philosophers who believe that we can know we're not being deceived by an evil demon (and so on) on purely *a priori* or non-perceptual grounds. These philosophers will deny premises (5) and (14).

The other philosophers who wish to resist the skeptic's conclusions have to deny SPK and SPJ. They will do so for a variety of different reasons. A **relevant alternatives theorist** denies that you have to "rule out," or have antecedent grounds for rejecting, *all* the propositions which are "bad" relative to what you purport to know by perception. In his view, you only have to have antecedent grounds for rejecting bad propositions *when those bad propositions are relevant*. Some **externalists** about justification say that, for your perceptual beliefs to be justified, those beliefs merely have to *be reliable*. You don't in addition have to *have evidence* that they're reliable, or that you're not being deceived by an evil demon, or anything of that sort. Some **coherentists** agree with the skeptic that the proposition *I am not being deceived* has to be part of your reasons for perceptual beliefs like *There are hands*. But on their view, propositions like *There are hands* may, in their turn, constitute part of your reason

for believing *I am not being deceived*. So neither proposition has to be justified *antecedently* to the other.

The **dogmatist** theories of perceptual knowledge and perceptual justification that I mentioned in §I provide another line of resistance to SPK and SPJ, different from all of the above. This is the response to the skeptic that I propose to explore.

I will concentrate on the skeptic about perceptual justification. This skeptic says that if you're to be justified in believing that things are as they perceptually seem to you, you need to have antecedent reason to believe that you're not in certain skeptical scenarios. The dogmatist about perceptual justification denies this. According to the dogmatist, when you have an experience as of *p*'s being the case, you have a kind of justification for believing *p* that does *not* presuppose or rest on any other evidence or justification you may have. You could have this justification even if there were *nothing* else you could appeal to as ampliative, non-question-begging evidence that *p* is the case. Hence, to be justified in believing *p*, you do *not* need to have the antecedent justification the skeptic demands. You do not need to have evidence that would enable you to rule the skeptic's scenarios out, in a non-question-begging way.

If this dogmatist account of perceptual justification is correct, then SPJ is false and we can resist the skeptic's argument. So we just need to determine whether dogmatism is correct.

III

I will now set out and defend a story about perceptual justification which says that dogmatism *is* correct.

To a first approximation, my view will be that whenever you have an experience as of *p*'s being the case, you thereby have immediate (*prima facie*) justification for believing *p*.

I need to explain what this means. I will begin with the notion of immediate justification. There are many extant criticisms of the "Myth of the Given" and of classical foundationalism. However, these criticisms still leave open the possibility of a respectable notion of "immediate justification."

Say that you are "mediately justified" in believing *p* iff you're justified in believing *p*, and this justification rests in part on the justification you have for believing other supporting propositions. Say that you are "immediately justified" in believing *p*, on the other hand, iff you're justified in believing *p*, and this justification doesn't rest on any evidence or justification you have for believing other propositions.

A few clarificatory remarks are in order.

First, the contrast between mediate and immediate justification has to do with the *source* of your justification, not the *strength* of your justification. We should not assume that immediately justified beliefs will be infallible or indubitable or anything like that. A belief might be fallible but immediately justi-

fied nonetheless. For instance, your knowledge of the reasons for which you act is fallible. (To take just one sort of mistake: you might have two motives for acting, and be wrong about which one you acted on.) Nonetheless, your beliefs about the reasons for which you acted are not ordinarily based on any evidence. So these beliefs are fallible; but in the ordinary case, they are immediately justified.

Second, the question whether your belief is mediately or immediately *justified* is not the same as the question how *psychologically* immediate or spontaneous your belief is. As we've already noted, your justification for believing one proposition can rest on your justification for believing another proposition, even when you do not believe the second proposition (but merely have justification for believing it), and hence, even when you do not infer the first proposition from the second. In the example I gave earlier, even though I do not actually believe that my gas gauge reads "E," I have *justification* for believing that, and my justification for believing that I'm out of gas rests on my justification for believing that the gas gauge reads "E." Hence, in that case my justification for believing that I'm out of gas is mediate. It would be mediate even if I were to form the belief that I'm out of gas *spontaneously*, upon seeing the gas gauge, without basing the belief that I'm out of gas on any beliefs about the gas gauge.

Third, the notion of immediate justification should not be confused with other epistemic notions in its vicinity. For instance, it should not be confused with the notion of a belief in a *self-evident* proposition, that is, a proposition such that anyone who understands it thereby has justification for believing it. (For example, the proposition that $1=1$ seems to be self-evident.) If there are self-evident propositions, then anyone who believes such a proposition will have immediate justification for believing it. But we should not assume that *only* beliefs in self-evident propositions can be immediately justified. That would need arguing. The notions of a self-evident belief and the notion of an immediately justified belief are different.

Nor is the notion of an immediately justified belief the same as the notion of a *self-justifying* belief, that is, a belief such that the mere fact that one holds the belief suffices to justify one in holding it. If there are self-justifying beliefs, then one's justification for believing them may be immediate. But we should not assume that *only* self-justifying beliefs can be immediately justified. The fact that a certain belief is immediately justified may be a *temporary* fact about the belief. You might be immediately justified in holding it at one time, but unjustified in holding it at another. A self-justifying belief, on the other hand, would have to be justified *whenever* one held it.

Nor is the notion of an immediately justified belief the same as the notion of an epistemically *autonomous* belief, that is, a belief which one could be justified in holding *without needing to hold any other beliefs*. Immediately justified beliefs don't have to be autonomous in that way. You might require certain background beliefs (perhaps even *justified* background beliefs) merely to

be able to entertain some belief B. That doesn't by itself show that your justification for believing B rests on your justification for those background beliefs. Compare the notion of *a priori* knowledge. When we ask whether a certain belief counts as an instance of *a priori* knowledge, we're not concerned with whether the subject *acquired the concepts* necessary to entertain that belief through experience. We're only concerned with the source of the subject's *justification* for the belief. What's necessary for him to entertain the belief is one matter, and what the nature of his justification for the belief is is another matter. In the same way, certain background beliefs might be required for you to *possess the concepts* necessary to entertain the belief B. That would show that B is not epistemically autonomous. But it does not, by itself, show that your justification for believing B rests on your justification for those background beliefs. So your justification for believing B can be immediate, even if the belief is not autonomous.²³

Fourth, when I say our senses give us immediate justification, the justification I have in mind is *prima facie* justification. *Prima facie* justification can be defeated or undermined by additional evidence. But in the absence of any such defeating evidence, *prima facie* justification for believing *p* will constitute all things considered justification for believing *p*.²⁴

Different sorts of things count as defeating evidence. My perceptual justification for believing *p* can be defeated by evidence in favor of not-*p*, by evidence that *p*'s truth is in these circumstances not ascertainable by perception, by evidence that my senses are malfunctioning, or by evidence that "explains away" its seeming to me that *p* is the case. The differences between these different sorts of defeaters will not matter here. However, I want us to understand "*prima facie*" and "defeating evidence" in such a way that only ordinary evidence of the sort employed by the man in the street and by the working scientist counts as defeating your *prima facie* justification. *A priori* skeptical arguments do not standardly introduce defeating evidence of that ordinary sort.²⁵ So I don't want us to talk like this:

The skeptic grants that our experiences give us *prima facie* justification for our perceptual beliefs, but if his philosophical arguments are sound, they defeat that justification.

Rather, if we use "*prima facie*" and "defeating evidence" in the way I propose, we ought to say this, instead:

The skeptic grants that our experiences purport or pre-theoretically seem to give us justification for our perceptual beliefs, but if his philosophical arguments are sound, they show that this is all an illusion. We do not have any justification (even *prima facie* justification) for beliefs about the external world, after all.²⁶

On my usage, if you have *prima facie* justification for believing something, and you have no (ordinary) evidence that defeats or undermines that *prima facie* justification, then you thereby have all things considered justification for your belief. I don't claim to be tracking ordinary usage perfectly here. This is a partly stipulative use of "*prima facie* justification."

Fifth, because the immediate justification we're talking about is merely *prima facie* justification, it's possible to have immediate justification for believing *p* and *also* to have mediate justification for believing *p*. Your beliefs can be "evidentially overdetermined."

Because of this, the claim that your perceptual beliefs are immediately justified need not commit one to all the other trappings of foundationalism. If you want, you can deny that it's possible to trace every justified belief's justification back to some foundational belief on which it rests. Considerations of coherence might sometimes, by themselves, suffice to justify beliefs. And perhaps all of your perceptual beliefs are justified in part by such considerations of coherence. That's consistent with the view I'm defending. I'm only claiming that your perceptual beliefs have *some* immediate justification—*some* justification which does not rest on your justification for believing further things.²⁷

Sixth, I do not suppose that every time you have a justified belief that *p* is the case, you will also have justification for believing that your belief *is* justified, and know what that justification is. What makes you justified in believing *p* is one thing; what makes you justified in believing *you have justification* for believing *p* is something else, something more sophisticated. This assumption is central to my approach. I will be giving an account of your justification for believing *the propositions that your perceptual experiences represent as being the case*. For example, what justifies you in believing that *there is a hand*, when your experiences represent that there is a hand? I am not offering an account of what justifies you in believing any *epistemic* propositions, like the proposition (i) that *you have justification* for believing that there is a hand; or the proposition (ii) that *your experiences give you justification* for believing that there is a hand. Nor am I here offering an account of what justifies you in believing (iii) that your perceptual faculties are *generally* reliable or trustworthy. I think your experiences can give you justification for believing that there is a hand, even if you lack justification for believing any of these more sophisticated propositions.²⁸

Seventh, I do not suppose that whenever you have justification for believing *p*, you will always be able to offer *reasons* or a *justifying argument* in support of your belief. It is important to distinguish between (i) the epistemic *status* of being justified, or having justification for believing something; and (ii) the *activity* of defending or giving a justifying argument for a claim. In my view, the status is epistemologically primary, and does not depend on your being able to engage in the activity. It can be reasonable for you to believe something even if you're not able to show that it's reasonable or explain what makes it reasonable. As Robert Audi writes:

It would seem that just as a little child can be of good character even if unable to defend its character against attack, one can have a justified belief even if, in response to someone who doubts this, one could not show that one does.²⁹

My view is that whenever you have an experience as of *p*, you thereby have immediate *prima facie* justification for believing *p*. (Perhaps you *also* have a kind of immediate introspective justification for believing *you have an experience as of p*; that's a question which I do not take up here.) Your experiences do not, in the same way, give you immediate *prima facie* justification for believing that you are dreaming, or being deceived by an evil demon, or that any of the skeptic's other hypotheses obtain.

Now we can understand what my view says. The next question is, should we believe it?

For a large class of propositions, like the proposition that there are hands, it's intuitively very natural to think that having an experience as of that proposition justifies one in believing that proposition to be true. What's more, one's justification here doesn't seem to depend on any complicated justifying argument. An experience as of there being hands seems to justify one in believing there are hands in a perfectly straightforward and immediate way. When asked, "What justifies you in believing there are hands?" one is likely to respond, "I can *simply see* that there are hands." One might be wrong: one might not really be seeing a hand. But it seems like *the mere fact* that one has a visual experience of that phenomenological sort is enough to make it reasonable for one to believe that there are hands. No *premises* about the character of one's experience—or any other sophisticated assumptions—seem to be needed.³⁰

I say, let's take these intuitive appearances at face value. Let's say that our perceptual beliefs in these propositions are indeed justified in a way that does not require any further beliefs or reflection or introspective awareness. They have a kind of justification which is *immediate*, albeit defeasible.³¹

One *might* try to buttress the claim that our experiences justify our perceptual beliefs by appeal to the fact that our experiences are reliably veridical. I will not do that. I do not want to base my arguments on considerations about reliability. (In fact, I believe that our perceptual justification would be in place no matter *how reliable* our experiences were. But I will not attempt to establish that here.³²)

It's not plausible that *every* proposition we believe on the basis of perception is immediately justified. I may believe on the basis of perception that a certain car is a Honda Accord; but this belief does not seem to be justified *just* by the fact that I have the experiences I have. So we have a challenge: we have to determine which of our perceptual beliefs are immediately justified, and which depend for their justification on *more* than just our having the experiences we have. We'll return to this challenge shortly.

We should not let that difficulty obscure the fact that our experiences *do* intuitively seem to justify us in believing a great many things about the external

world, all by themselves. Some examples: there is a light ahead; there is something wet and noisy ahead (a belief you might form while standing on the beach with your eyes closed); there is something solid here (a belief you form while pressing against a wall). For these propositions and many others like them, it's very tempting to say that we have immediate justification for believing them, justification which does not derive from any evidence or justification we have for believing further propositions.

Philosophers have only passed over this natural view because they've thought it vulnerable to various objections. However, that judgement was too hasty. The resources of this natural view have been sorely underestimated.

For instance, many philosophers have assumed that our perceptual beliefs can't be immediately justified, because they're fallible, and they can be defeated or undermined by additional evidence. Others have argued that our perceptual beliefs can't be immediately justified because they're not self-evident, or self-justifying, or epistemically autonomous. But as I've explained the notion of immediate justification, it's quite compatible with all of that. Once it's recognized how modest the claim that our perceptual beliefs are immediately justified really is, many of these standard objections simply evaporate.

Though the claim that our perceptual beliefs are immediately justified is modest, it still has anti-skeptical punch. For this view is a species of dogmatism about perceptual justification, and as we've seen, a case for dogmatism is a case against the skeptic's principle SPJ. According to the dogmatist, you can be justified in believing that there is a hand, even if you're not in a position to offer any non-question-begging argument for that proposition—in other words, even if you're not able to defend that proposition by appeal to premises you have antecedent justification for believing. In particular, you can be justified in believing that there is a hand, even if you have no antecedent justification for believing that you're not being deceived by an evil demon. It should be clear that if our perceptual beliefs are immediately justified, then the dogmatist is right. If our perceptual beliefs are immediately justified, then they are justified in a way that doesn't rely on any other antecedently justified beliefs whatsoever. One who has an immediately justified belief need have nothing he could appeal to by way of non-question-begging argument in support of the belief. He is justified for all that.³³

Hence if we accept my account of perceptual justification, then we can reject the skeptic's key principle, SPJ. We'll then have succeeded at the modest anti-skeptical project. We'll have shown how we can retain our conviction that we have perceptual justification, without *also* accepting the premises the skeptic needs to show we *can't* have such justification.

I don't want to claim that you *never* have to rule out skeptical hypotheses. I claim merely that your experiences give you *prima facie* justification for your perceptual beliefs, and that it's not a precondition of having this *prima facie* justification that you be able to rule out any skeptical hypotheses. This *prima facie* justification can be undermined or threatened if you gain positive empir-

ical evidence that *you really are in a skeptical scenario*. (For instance, if a ticker tape appears at the bottom of your visual field with the words “You are a brain in a vat...”³⁴) If you acquire evidence of that sort, then you’d have to find some non-question-begging way of ruling the skeptical hypothesis out, before you’d be *all things considered justified* in believing that things are as your experiences present them.³⁵ In the standard case, though, when the *prima facie* justification you get from your experiences is not defeated or undermined, then it counts as all things considered justification, without your having to do this.³⁶

Note that I am only proposing a dogmatic story *about the nature* of your justification. *My argument* that your justification has this nature is not itself dogmatic. The argument that your justification has this nature proceeds via standard philosophical methodology: we start with what it seems intuitively natural to say about perception, and we retain that natural view until we find objections that require us to abandon it. This is just sensible philosophical conservatism.³⁷

Let’s return to the question: What sorts of beliefs do we have this immediate perceptual justification for?

Some of the beliefs we form by perception are clearly more epistemologically sophisticated than others. I gave the example earlier of looking at my gas gauge, seeing that it reads “E,” and forming the belief that I’m out of gas. This belief was formed by perception, but it seems to go beyond the strict deliverances of my perceptual experiences. My justification for believing the car is out of gas isn’t wholly perceptual. It also rests on background knowledge I have about cars and gas gauges. It’s true that I needn’t have *inferred* the belief that I’m out of gas from any other premises. I might have formed it as a natural and spontaneous result of seeing the gas gauge. Nonetheless, my justification for the belief that I’m out of gas seems to draw on more than just what’s provided by my visual experiences. Contrast the case where I have an experience as of there being a light ahead, and form the belief that there is a light ahead. Here it seems I do have justification for my belief which is wholly perceptual. (I might *also* have some not-wholly-perceptual justification, as well. Perhaps someone told me that there was a light ahead. But my perceptual justification could be in place even if this additional justification were lacking.)

Let’s distinguish among the beliefs we form by perception. Some propositions are such that we see or seem to see that they are so in virtue of seeming to see that other propositions are so. For instance, I seem to see that there’s a policeman ahead partly in virtue of seeming to see that there’s a blue-coated figure ahead, and partly in virtue of having certain background evidence about the ways that members of our society typically dress. Perhaps we ought strictly speaking to deny that my perceptual experiences themselves have the content: *There is a policeman ahead*. Much of our use of locutions like “It looks as if...” and “I (seem to) see that...” is influenced not just by what representational contents our experiences have, but also by what further conclusions we take those experiences to make obvious. For instance, when you look at someone’s face

as he comes out of an examination and, as you'd put it, *see that he passed the test*, the proposition that he passed the test is not itself represented by any of your perceptual experiences. Rather, you only see that he passed the test in virtue of seeing that he has a certain kind of expression on his face. I will call those propositions we seem to perceive to be so, but *not* in virtue of seeming to perceive that other propositions are so, **perceptually basic** propositions, or propositions that our experiences **basically represent**. (I will not dwell on the question whether non-basic propositions ever *really* deserve to be called propositions we "seem to see" to be the case.) The propositions that there is a policeman ahead, and that a certain person passed a test, are not perceptually basic. I think it is only perceptually basic propositions which purport to be justified just by the deliverances of our current perceptual experiences.

Note that perceptually basic propositions are *not* about sense-data or the character of our experiences. They are about manifest observable properties of objects in the world. Those are the sorts of things that our experiences represent to us. And it need not be the case that perceptually basic propositions only concern the facing surfaces of objects. The cognitive psychologists will have to tell us whether our experiences have contents like *There is a (complete) hand*, or whether they instead have contents like *There is a facing flesh-colored surface of such-and-such a shape*.

It is likely that what's perceptually basic will vary for different people, and it may also vary for the same person over time. (Certain kinds of training or aging might cause you to represent some phenomena by sight which your experiences did not previously represent. These phenomena may turn out to be *conceptually* quite sophisticated. For instance, it may be possible to have visual experiences which basically represent a given painting as being expressionist.) I take no official stand on how far the class of basically represented propositions extends; and I think it is difficult to say in any particular case exactly which propositions are perceptually basic. This is because it is very difficult to distinguish between the content of one's experiences, and the contents of the beliefs that one forms as a result of having those experiences. Fortunately, for our present purposes, we only need a schematic grasp of the notion of a perceptually basic proposition.

The official version of my view is that we have immediate *prima facie* justification for believing those propositions that our experiences basically represent to us—whichever propositions those turn out to be.³⁸ Perhaps we don't *believe* very many perceptually basic propositions. However, our experiences *do* basically represent some propositions to us—that is, they represent some propositions to us not by virtue of representing other propositions to us—and I claim that when our experiences do this, we thereby acquire immediate justification for believing that those basically represented propositions are the case. Perhaps our minds pass over them to more sophisticated propositions, as in my example with the gas gauge. My claim is merely that our immediate justification for the perceptually basic propositions is playing a crucial role in the epistemological story.

IV

I want briefly to respond to an important objection to my account of perceptual justification. That objection goes: “All observation is theory-laden. Hence none of your perceptual beliefs can be justified *just* by your having the experiences you do. Your justification for any perceptual belief always derives in part from the background theory you hold.”

There are a number of *different senses* in which observation might be “theory-laden,” and these are not always clearly distinguished. Several of these pose no obstacle to the view I’m defending. So it’s important to be sure whether our grounds for saying that all observation is theory-laden are grounds for saying that observation is *theory-laden in a sense that prevents our perceptual beliefs from being immediately justified*.

The claim “observation is theory-laden” might mean that what theory you hold can causally affect what experiences you have. Two subjects given the same sensory stimulation but with different background beliefs can differ with respect to what perceptually seems to them to be the case. There is a substantial amount of empirical evidence that this is so. For instance, if you believe that the object you’re looking at is a familiar object like a carrot, you’re likely to experience it as being more orange than you would if you lacked that belief.

Does this pose any obstacle to the view I’m defending? Does it show that your justification for believing that the object is orange cannot be immediate?

It does not. I’m concerned with which transitions *from experience to belief* would result in justified belief. The present claim concerns how one comes to have the experiences, in the first place. These are independent issues. Why should the fact that your background beliefs causally affect what experiences you have show that the justification you get from those experiences relies on or derives from those background beliefs? Your sunglasses causally affect your experiences, but none of your perceptual beliefs are justified to any extent by your sunglasses. So if your background beliefs are playing a role analogous to your sunglasses, why should that fact alone make them contribute to the justification of your perceptual belief?³⁹

The claim “observation is theory-laden” might instead mean that background beliefs are capable of *defeating* any justification you get from your senses, and hence that your background beliefs play an indispensable role in determining what you’re all things considered justified in believing. For instance, if you have the background belief that the present lighting conditions systematically make things look redder than they actually are, then what you’re all things considered justified in believing on the basis of a given course of visual experience will be different from what it would be if you lacked this background belief.

I have no argument with this. On my view, the immediate justification you get from your experiences is only *prima facie* justification. I’m quite happy to say that your background beliefs play a role in determining when that *prima facie* justification is defeated or undermined, and when it isn’t.

The claim “observation is theory-laden” might mean that you need to have certain theoretical beliefs before you’re even able to *entertain* or *form* certain observational beliefs. (This is the view we express when we say that your descriptions of the observational evidence are “contaminated by” concepts that you have only because you accept the theory.) For instance, in order to be able to entertain the proposition that the table is *solid*, you need the concept of solidity, and to have that, you may need to believe a certain “folk physics.” You may need to have beliefs about how solid objects interact with other things.

Again, if this is true, I have no argument with it. I claim that your perceptual beliefs are immediately justified, not that they are autonomously justified. You may need certain background beliefs to be able to *entertain* your perceptual beliefs. I’m only claiming that your *justification* for your perceptual beliefs doesn’t rest on those background beliefs. As I said before, what’s necessary for you to entertain your perceptual beliefs is one matter, and where your justification for those beliefs comes from is another matter.

Finally, the claim “observation is theory-laden” might mean that background theory necessarily plays a role in your acquisition of even *prima facie* justification from your senses. This does contradict the view I’m defending. However, it’s quite unclear whether we have any good reasons to believe that observation is theory-laden in this sense. The standard arguments for believing that observation is theory-laden are usually aimed at establishing one of the earlier claims. My account of perceptual justification is compatible with the claim that “observation is theory-laden” in any of the senses of that claim in which it is uncontroversially true.

I have argued for a certain account of how our experiences as of *p* give us perceptual justification for believing *p*. Because the connections between justification and knowledge are subtle and complicated, I have not addressed the question how we acquire perceptual knowledge—though I believe that in the ordinary case, when the justification one gets from one’s experience is undefeated, one does have perceptual knowledge. And I have not provided any recipe for refuting the skeptic using only premises he’ll accept. What I have done is offer a plausible and intuitive account of perceptual justification that *we* can accept. I have also shown how, once we accept this account of perceptual justification, the skeptic’s best argument is revealed to rest on a false principle. Contrary to what the skeptic says, it is *not* a precondition for having perceptual justification for believing *p* that one have other, antecedently justified beliefs that one could appeal to without begging the question against the skeptic. What it’s reasonable for us to believe can outstrip what we’re able to defend in a non-question-begging way.⁴⁰

Notes

¹What follows is not a serious attempt at Moore exegesis. I’m just using Moore to introduce some ideas I want to develop. For a more serious examination of Moore’s own views, see the final two chapters of Baldwin 1990.

²Moore 1939, pp. 168–70; see also Moore 1909, pp. 159–60.

³To be more exact: any argument Moore could offer would either (i) appeal to further premises that the skeptic finds question-begging, or (ii) appeal to premises of whose truth Moore is far less certain than he is that he has a hand. For ease of exposition, I set this second possibility aside.

⁴This model is endorsed in Evans 1982 Ch. 5 and 7; Searle 1983 Ch. 2; Peacocke 1983 Ch. 1; Peacocke 1992 Ch. 3; Burge 1986; McGinn 1989 Ch. 1; Block 1990; Block forthcoming; Shoemaker 1990; and Shoemaker 1994 Lecture III.

Some philosophers go even further and claim that an experience's representational properties fully determine "what it's like" to have that experience. These philosophers are often called **intentionalists**. See for instance: Harman 1990; Tye 1992; Tye 1994; Tye 1995; Dretske 1995; Byrne and Hilbert 1997; Lycan 1996; Dennett 1988; and Dennett 1991. (Armstrong 1961 Ch. 9–10, Armstrong 1968 Ch. 10, and Pitcher 1971 Ch. 2 were early intentionalists, who tried to analyze experiences in terms of dispositions to believe. Contemporary intentionalists take experiences to be autonomous propositional attitudes, not analyzable in terms of beliefs or dispositions to believe.) I am sympathetic to intentionalism, but it is not required by the model I mean to be employing here.

⁵Plantinga emphasizes the difference between (i) the claim that what warrants your perceptual belief is your *believing* that you have certain kinds of experiences, and (ii) the claim that what warrants your perceptual belief is your simply *having* those experiences. See Plantinga 1993b, pp. 95ff and 183ff. See also the works cited in note 31, below.

⁶A typical fallibilist will argue against the skeptic with an abductive argument of this sort: (i) I have an experience as of there being a hand; the best explanation of why I have this experience entails that there is a hand; hence, there is a hand.

Someone who both accepts dogmatism and is aware of what experiences he's having may be in a position to argue against the skeptic as follows: (ii) I have an experience as of there being a hand, which (according to dogmatism) makes me justified in believing there is a hand; hence, there is a hand. However, the dogmatist does not think that our justification for our perceptual beliefs *ordinarily* proceeds via argument like (ii). This is for two reasons.

First, as I've emphasized, the dogmatist thinks we have a kind of justification for our perceptual beliefs that does not require us to appeal to awareness of our experiences, or to premises about what sorts of experiences we're having.

Second, an inference can give one justification for believing its conclusion only if one is entitled to rely on it. For many forms of inference, like *modus ponens* and simple inductive arguments, one can be entitled to rely on them without doing any philosophy. The dogmatist does not claim that this is one of those cases. So far as the dogmatist is concerned, it may be that argument (ii) gives one justification for believing its conclusion *only if one appreciates the philosophical merits of dogmatism*. This is in fact my own view. I believe that the ordinary person is not in a position to employ argument (ii), even if he is aware of his experiences. The argument relies on a step he has no right to assume is epistemologically legitimate, until he does philosophy. The ordinary person *does* have justification for *the conclusion* of argument (ii)—but that justification comes from his merely having certain experiences, not from his rehearsing argument (ii) to himself.

⁷For other examinations and diagnoses of the skeptic's reasoning, see Klein 1981; Stroud 1984; Wright 1991; Sosa 1988; Brueckner 1994; and Cohen 1998. The account of the skeptic's reasoning I offer is similar to one of the accounts Klein discusses.

⁸See Feldman and Conee 1985 on the contrast between "well-founded" and "ill-founded" beliefs. As I explain in Pryor forthcoming §3.2, we need such a contrast even when we're dealing with beliefs which are immediately justified. We just cannot take the time to explore it here.

⁹See, for instance, Stroud 1984 Ch. 1; Dancy 1985 Ch. 1; Wright 1985, p. 129; and Wright 1991.

Crispin Wright considers an argument against the possibility of justified perceptual belief in Wright 1991. (He thinks this argument can be successfully countered.) Wright's skeptical argument is not suitable for my purposes, since it employs an externalist notion of justification. (Wright terms this notion "warrant.") I want to discuss skeptical challenges to the possibility of our having any *internalist* grounds for believing that the world is as it appears to us. There are also several

problems with Wright's argument which I'd like to avoid. See Brueckner 1992 and Tymoczko and Vogel 1992 for discussion.

¹⁰Here the skeptic relies on the following modest Closure Principle: If you know *p*, and you know that *p* implies *q*, then either you already know *q* or you are in a position to deduce and thereby come to know *q*.

¹¹Some philosophers think it is obvious that we don't know we're not being deceived. See for instance Cohen 1988; DeRose 1995; and especially Nozick 1981, p. 201. To other philosophers, like Moore, it seems equally obvious that we *do* know we're not being deceived by an evil demon. We derive that knowledge in part from our perceptual knowledge that there are hands, tables, and the like all around us; which would not be the case if we were being deceived by an evil demon. See Moore 1959, pp. 190–91 and Klein 1981 §2.10.

The skeptic could try to *argue for* his premise that you can't know you're not being deceived by an evil demon, as follows:

- (i) It's metaphysically possible for someone to be deceived by an evil demon, compatibly with having all the experiences you're now having.
- (ii) Since knowledge is factive, someone who *is* being deceived by an evil demon can't know on the basis of his experiences that he's *not* being deceived by an evil demon.
- (iii) Since you have the same experiences he has, for all you know on the basis of your experiences, you may right now be in his situation.
- (iv) So you're in no better position to know you're not being deceived by an evil demon than he is.

But a clear-sighted fallibilist who recognized the difference between metaphysical possibility and epistemic possibility would not find this reasoning to have any force. The argument moves from considerations about what's metaphysically possible to a conclusion about what's left epistemically open. No fallibilist ought to accept that move. The fallibilist's central thesis is precisely that it's possible to know things on the basis of defeasible evidence, evidence the possession of which is metaphysically compatible with your being wrong. According to the fallibilist, the fact that some evidence is metaphysically compatible with *q*'s being false is *not enough* to show that you can't know *q* on the basis of that evidence. (This point was made by Moore in Moore 1959, see esp. pp. 176–77. The difference between epistemological and metaphysical possibility is also one of the lessons of Kripke 1972, and the literature that followed it.)

Of course, the fact that we accept a fallibilist conception of knowledge will not give us the means to *refute* a skeptic who advances the argument (i)–(iv). We're not likely to persuade the skeptic to convert to fallibilism. But if our aim is merely the modest anti-skeptical project, that shouldn't concern us. The modest anti-skeptical project just aims to establish to *our* satisfaction that we can have perceptual knowledge, without contradicting any of the other things we accept about perception and about what would be required for perceptual knowledge. Since the argument (i)–(iv) relies on a move that no fallibilist should be inclined to grant, it poses no threat to our completing the modest anti-skeptical project.

A further problem with the argument (i)–(iv) is that it cannot easily be translated into an argument that would be usable by skeptics who challenge our possession of justification rather than knowledge (see the discussion which follows in the text).

¹²For this reason, Stroud 1984 Ch. 1 appeals not to any simple Closure Principle, but rather to the following principle: "If you know that *q* is incompatible with *your knowing p*, then you know *p* only if you also know not-*q*." On very modest assumptions, Stroud's Principle is equivalent to the Closure Principle "If you know that *q* is incompatible with *p*, then you know *p* only if you also know not-*q*" plus the KK Principle "If you know *p*, then you know that you know *p*."

I selected the evil demon hypothesis to make our discussion simpler. However, as we proceed, we will discover that the skeptic needs to appeal to a principle different in spirit from any Closure Principle. In the long run it will not matter which skeptical hypothesis we employ.

¹³But the story about *how* our knowledge that we're not being deceived rests on our perceptual knowledge about our environment is complicated. See note 35, below.

¹⁴See BonJour 1985 Ch. 8, and "inference to the best explanation" replies to the skeptic, of the sort offered by Russell, Harman 1973 Ch. 8 and 11, and Vogel 1990b.

¹⁵And even if my suspicions were warranted, it would still seem plausible to say that I have *prima facie* justification for believing that I'm out of gas. It's just that I'd *also* have reason to believe that I'm hallucinating, and this would *defeat* my *prima facie* justification for believing that I'm out of gas. More on *prima facie* justification in §III, below.

¹⁶This notion of antecedent justification is also germane to discussions of relevant alternatives responses to skepticism. See §1.2 of Pryor forthcoming.

¹⁷By Strengthened Closure, knowing that $I=I$ would require you to be in a position to *antecedently* know that $snow\ is\ white \supset I=I$; hence your knowledge that $snow\ is\ white \supset I=I$ could not be *based on* your knowledge that $I=I$. You would have to have independent reasons for believing that $snow\ is\ white \supset I=I$, if you're to count as knowing that $I=I$.

¹⁸DeRose 1995 treats as "bad" those skeptical hypotheses *h* which are such that if they obtained, we'd still believe *p* but our beliefs would be false, and *h* gives an explanation of why we'd still believe *p* in that situation. Cohen 1988 and Cohen 1999 treat as "bad" those skeptical hypotheses which entail that we have all the evidence for *p* that we actually have. I take these to be variants on the same basic theme.

¹⁹On the notion of undermining, see Pollock 1974 Ch. 2 and 5; Pollock 1984; Pollock and Cruz 1999, pp. 195–97; Cohen 1987; Alston 1989, p. 238; Wright 1982, pp. 117–20; Wright 1986, pp. 290–91; Wright 1991, pp. 94–95; Goldman 1986 Ch. 4; and Audi 1993, pp. 142–44. These authors use a variety of different terms for this notion.

²⁰It may be that the skeptic is unable to articulate any satisfactory *general* principle about what makes a possibility a "bad" one, but that we nonetheless have intuitions about *particular* skeptical possibilities, like the possibilities that one is dreaming or being deceived by a demon. I have no real objection to this. I do not insist that the skeptic derive (7) from general principles. I only insist that he give us *some* sort of story about why the demand to antecedently rule out alternatives is imposed *when we're dealing with the alternatives he mentions*. (We have seen no reason to impose this demand across the board, for *every* alternative to what we purport to know.)

²¹See Stroud 1984, pp. 18–23, and 30–31.

²²Stroud does not himself acknowledge this. He repeatedly calls the skeptic's condition a "necessary condition" for perceptual knowledge, and never calls it a precondition. Nor does he make any explicit appeal to considerations of epistemic priority. See Sosa 1988 pp. 158–59 for more on these features of Stroud's discussion.

Considerations of epistemic priority also play an important role in Klein's and Wright's discussions: see Klein 1981 §§2.13–15; Klein 1995, n. 16; Wright 1985, pp. 433 and 435–38; Wright 1991, pp. 99–100; and Brueckner 1992, pp. 313–15.

²³See Audi 1993 Ch. 3; Alston 1989, pp. 62–63, 293; Burge 1993, p. 460; and Plantinga 1993a, p. 71fn.

²⁴It will constitute all things considered justification but it need not always constitute *very much* justification. All things considered, you might be more justified in believing *p* than *not-p*, though the balance of evidence you possess in favor of *p* is *not very great*—not enough to enable you to *know p*, for instance. This raises further subtleties about the connection between justification and knowledge that we can not explore here. I believe that in many cases we *do* get enough *prima facie* justification from our experiences to have perceptual knowledge.

²⁵Though see note 34, below.

²⁶We might have been *blameless* in accepting those beliefs, until we heard the skeptic's arguments. But the relations between blamelessness and epistemic justification are not straightforward. See Pryor forthcoming for discussion.

²⁷As I mention in note 24, I do believe that our perceptual beliefs give us *enough* justification, all by themselves, to sometimes constitute knowledge. But I will not defend that claim here.

²⁸BonJour at one time argued that to be justified in believing *p*, you do need justification for believing that your belief has some feature in virtue of which it is justified. See BonJour 1978 pp. 5–6, and BonJour 1985 §2.3. This view seems to me to confuse justification with a more sophisticated and harder-to-achieve epistemic standing. It is a particularly strong form of the view I call “Access Internalism” in Pryor forthcoming.

Some philosophers think that you can be adequately justified in believing *p* only if you’re also justified in believing that *there are no defeaters* of your justification for believing *p*. (I encounter this claim more often in conversation than in print.) This view also seems implausible to me. In the first place, our primary target here is what it takes to have *prima facie* justification for believing *p*. And the claim that you have *prima facie* justification is compatible with the claim that your *prima facie* justification is defeated by other evidence you possess. So, since *prima facie* justification is compatible with the existence of defeaters, why should it be a requirement for having *prima facie* justification that one be able to rule out the existence of defeaters? At best, it would only be *all things considered* justification for believing *p* which required one to rule out the existence of defeaters.

Even this claim, though—the claim that *all things considered* justification for believing *p* requires one to rule out the existence of defeaters—seems doubtful to me. It’s important to distinguish two claims: (i) If one acquires defeating evidence against *p*, then in order to be all things considered justified in believing *p*, one has to have reasons that outweigh or undermine that defeating evidence. This is surely correct. (ii) In order to be all things considered justified in believing *p*, one has to be justified in believing that one’s justification for believing *p* is undefeated. This is what I doubt. In my view, having justification for believing *p* does not require you to have justification for believing that you have *any* justification for believing *p*. So why should it require you to have justification for believing that you have *undefeated* justification for believing *p*?

²⁹Audi 1993 Ch. 4, at p. 145. See also Audi 1993 Ch. 10 and 12; Klein 1981, pp. 6–9; Alston 1989 Essays 1–3; and Alston 1999.

³⁰I am interpreting a bit here. When one says “I can *simply see* that there are hands,” I assume that one is reporting *the facts* that make one justified, not *some premise* on which one’s justification rests. I also assume that whether or not the experience counts as a case of *veridical seeing* is irrelevant to one’s justification. Both of these assumptions can be and have been disputed. But they are natural assumptions to make.

³¹Several other authors have also claimed that our perceptual beliefs can be immediately justified. See for instance Alston 1989 Essays 1–3; Pollock 1974 Ch. 2–5; Pollock and Cruz 1999 Ch. 5 §4.2 and Ch. 7 §§1–3; and Audi 1993 esp. Ch. 3, 4, 10, and 12. On some interpretations of Chisholm’s views, our perceptual beliefs come out immediately justified; see Alston 1997 for a careful discussion of this.

³²See Pryor forthcoming for a discussion of different kinds of internalism, and their motivations.

³³Strictly speaking, the claim that one’s perceptual beliefs are immediately justified does not entail that SPJ is false. It only undercuts the most obvious line of reasoning in support of SPJ. That line of reasoning says that one needs antecedent reasons for believing that one is not in the skeptic’s scenarios because one’s justification for one’s perceptual beliefs *has to rest* in part on one’s antecedent right to believe that one is not in the skeptic’s scenarios.

It would be open to the skeptic to forego this line of reasoning, and attempt to support SPJ by some other means. He would have to argue that having antecedent reasons for believing that one is not in the skeptic’s scenarios is a mere *necessary condition* for having perceptual justification, and not something on which one’s perceptual justification *rests*. It is quite unclear how the skeptic could motivate this claim. In particular, it is quite unclear how he could justify the insistence that one’s reasons to believe that one is not in the skeptical scenarios *be antecedent to* any of one’s perceptual justification, when he has granted that one’s perceptual justification does not *rest on* those reasons.

³⁴More problematic are cases where you get statistical evidence, such as evidence that 7 out of 10 experiencing subjects are brains in vats, or evidence that *you* are a brain in a vat 7 mornings out of 10. I think that defeaters of this statistical sort threaten knowledge in ways which are very dif-

ferent from the skeptical arguments we've been examining so far. We cannot explore these issues here. For some discussion, see Vogel's remarks on "Car-Theft Cases," in Vogel 1990a.

³⁵I do not have the space here to discuss how we'd *go about* ruling out a skeptical hypothesis, either in the case where we have positive evidence that it obtains, or in the ordinary case, where we have no evidence in favor of it but we're just curious whether or not it obtains. But a few brief remarks may help clarify my position.

Suppose U is some hypothesis such that the only reasons you have for believing U to be false presuppose the truth of *p*, but if you were to acquire evidence for U, that would defeat or undermine your justification for believing *p*. (For example, let *p* be some body of perceptual beliefs, and let U be a belief like *I am a brain in a vat* or *My senses are unreliable*. Alternatively, let *p* be *My friend will meet me tomorrow in NY as she promised* and let U be *My friend will be in LA tomorrow, as a result of being kidnapped and held hostage*.)

There are three accounts one can give of such cases. According to the first account, although you are justified in believing *p*, and *some* arguments from *p* seem like they could justify belief in not-U, the particular kinds of grounds you have for *p* do not enable *p* to provide you with justification for believing not-U. (For example, if God told you that *p* is true, you might be able to justifiably infer not-U from *p*, but when your grounds for believing *p* are perceptual, you can't justifiably infer not-U from *p*.) This account allows that you could still justifiably believe *p* on the grounds in question.

According to the second account, not only does the justification you have for believing *p* fail to justify you in believing not-U: in addition, your justification for believing *p* somehow *presupposes* that U is false. So you wouldn't be justified in believing *p* unless you were also justified in believing not-U, and since your justification for believing *p* can't *provide* any justification for believing not-U, you need to have independent reason to believe not-U if you're to be justified in believing *p*. In such cases, principles like the skeptic's SPK and SPJ seem to be genuinely applicable.

According to the third account, your justification for believing *p* *does* give you justification for believing not-U. However, because U is a potential defeater or underminer of your justification for believing *p*, any evidence you acquired in favor of U would defeat (or at least contribute towards the defeat of) your justification for believing *p*.

In my view, each of these accounts is correct for some cases. See Wright 1985; Davies 1998; Davies forthcoming; Wright forthcoming; Feldman 1995; and Sorensen 1988 for discussion of some of the relevant cases. The interesting question is: What account is correct when *p* is something we purport to know on the basis of perception and U is a "bad" skeptical hypothesis?

My arguments in this paper say that the second account is wrong for those cases. My inclination is to treat those cases as all falling under the third account. However, I will have to defend that inclination on another occasion. One relevant question is whether your perceptual experiences only give you justification for believing simple propositions like *There is a hand*, or whether, if you are sufficiently reflective and conceptually well-equipped, they also give you justification for believing the more sophisticated proposition that *you see a hand*. Although I have only discussed the simpler case here, I believe your perceptual experiences give you justification for believing both sorts of proposition. (Your justification for believing that there is a hand doesn't *rest* on your justification for believing that you see a hand, though, any more than it rests on your justification for believing that you have certain sorts of experiences. Your justification for these beliefs just comes from the same source.)

³⁶Suppose (i) one were to adopt a relevant alternatives account of perceptual *justification*, and (ii) one counted an alternative as relevant iff the subject has some positive *prima facie* evidence that it obtains, and (iii) one said that "ruling out" an alternative consisted in having antecedent justification for believing—contrary to the *prima facie* evidence—that the alternative does *not* obtain. I don't know of anyone who holds this view, but it's certainly a possible view. Such a relevant alternatives theorist would agree with me about when we have all things considered justification for our perceptual beliefs, and about what's required to have that justification. Hence, he will also reject the skeptic's principle SPJ. The arguments for his view may be different than the arguments

for mine. (He says nothing about immediate justification or *prima facie* justification; I say nothing about “relevance” or “irrelevance.”) Then again, it may turn out that we have the same view, in different clothing.

I would disagree with a relevant alternatives theorist of justification who said that alternatives can sometimes be relevant, and need to be ruled out, even when the subject lacks positive evidence that they obtain. When we lack such positive evidence, I think that our justification for our perceptual beliefs does not require us to have antecedent justification for believing *anything* else, no matter how “relevant.”

³⁷In my view, relying on this philosophical conservatism is a perfectly legitimate and satisfactory way to engage in the modest anti-skeptical project. However, if we could get one, we would certainly *like* to have a more informative story about why our perceptual experiences offer us the justification they do.

One such story would appeal to the fact that our perceptual beliefs are *irresistible*. This story would say things like this: “The claim that I *ought* to believe *p* implies that I *can* believe *p*. Now, if my belief in *p* is irresistible—if the only thing I *can* believe in a certain situation is *p*—then what I *ought* to believe in that situation must therefore be *p*.” A different story would make it constitutive of our concept of justification, or of our perceptual concepts, that having such-and-such experiences counts as good reason to believe that certain perceptible properties are instantiated.

I am not attracted to any of those stories. In my view, it’s not the irresistibility of our perceptual beliefs, nor the nature of our concepts, which explains why our experiences give us the immediate justification they do. Rather, it’s the peculiar “phenomenal force” or way our experiences have of presenting propositions to us. Our experience represent propositions in such a way that it “feels as if” we could tell that those propositions are true—and that we’re perceiving them to be true—just by virtue of having them so represented. (Of course, to be able to articulate this “feeling” takes a high grade of reflective awareness.) I think this “feeling” is part of what distinguishes the attitude of experiencing that *p* from other propositional attitudes, like belief and visual imagination. Beliefs and visual images might come to us irresistibly, without having that kind of “phenomenal force.” (Perhaps that’s what happens in cases of super-blindsight.) It is difficult to explain what this “phenomenal force” amounts to, but I think that it is an important notion, and that it needs to be part of the story about why our experiences give us the justification they do. I will have to develop these suggestive remarks elsewhere.

³⁸Notice: I am *only arguing* that we have immediate justification for believing what’s basically represented. I do not argue that we *only have* immediate justification for believing what’s basically represented. Perhaps we have immediate justification for believing some of the things that our experiences *non*-basically “represent,” too. (I doubt that, but I am not going to defend these doubts here.)

³⁹Worry: But if I have different background beliefs than you, then we might be subjected to the same stimulus, and yet end up with different experiences, and so different perceptual beliefs. Can *both* of our beliefs be justified? Answer: Why not? I see no reason to rule out the possibility that one and the same causal stimulus will result in people having justification for incompatible beliefs. If I have different eyeballs than you, then we might be subjected to the same stimulus, and yet end up with different experiences, and so different perceptual beliefs. No one thinks *that’s* an obstacle to our both having *prima facie* justification for our respective beliefs. So why should matters be any different when it’s our background beliefs which are affecting what experiences we have?

⁴⁰I owe thanks to many people for their advice and feedback about this paper and its ancestors: Alex Byrne, Andrew Chignell, Stewart Cohen, Joe Cruz, Christoph Erlenkamp, Steve Gross, Ned Hall, Aaron James, Mark Johnston, Mark Kalderon, Tom Kelly, Adam Leite, Dave Macarthur, Tim Maudlin, Elijah Millgram, Dick Moran, Jennifer Noonan, Christian Piller, Ian Proops, Gideon Rosen, Susanna Siegel, Alison Simmons, Scott Soames, Mike Thau, and Ralph Wedgwood. My warmest thanks to all of them. Thanks also to audiences at York University, Columbia University, and NYU.

References

- William Alston, *Epistemic Justification* (Cornell Univ. Press, Ithaca, 1989)
- William Alston, "Chisholm on the epistemology of perception" in Lewis E. Hahn, ed., *The Philosophy of Roderick Chisholm* (Open Court, La Salle, Ill., 1997), 107–25
- William Alston, "Perceptual knowledge" in John Greco and Ernest Sosa, eds., *The Blackwell Guide to Epistemology* (Blackwell, Malden, Mass., 1999), 223–42
- David Armstrong, *Perception and the Physical World* (Routledge and Kegan Paul, London, 1961)
- David Armstrong, *A Materialist Theory of the Mind* (Humanities Press, New York, 1968)
- Robert Audi, *The Structure of Justification* (Cambridge Univ. Press, Cambridge, 1993)
- Thomas Baldwin, *G.E. Moore* (Routledge and Kegan Paul, London, 1990)
- Thomas Baldwin, ed., *G.E. Moore: Selected Writings* (Routledge and Kegan Paul, New York, 1993)
- Ned Block, "Inverted earth" in James Tomberlin, ed., *Philosophical Perspectives* 4 (1990), 53–79
- Ned Block, "Mental paint," forthcoming; available at <<http://www.nyu.edu/gsas/dept/philo/faculty/block/>>
- Laurence Bonjour, "Can empirical knowledge have a foundation?" *American Philosophical Quarterly* 15 (1978), 1–13
- Laurence Bonjour, *The Structure of Empirical Knowledge* (Harvard Univ. Press, Cambridge, Mass., 1985)
- Anthony Brueckner, "Problems with the Wright route to skepticism," *Mind* 101 (1992), 309–17
- Anthony Brueckner, "The structure of the skeptical argument," *Philosophy and Phenomenological Research* 54 (1994), 827–35
- Tyler Burge, "Cartesian error and the objectivity of perception" in Philip Pettit and John McDowell, eds., *Subject, Thought and Context* (Oxford Univ. Press, Oxford, 1986), 117–36
- Tyler Burge, "Content preservation," *Philosophical Review* 102 (1993), 457–88
- Alex Byrne and David R. Hilbert, "Colors and reflectances" in Alex Byrne and David R. Hilbert, eds., *Readings on Color*, vol. 1 (MIT Press, Cambridge, Mass., 1997), 263–88
- Stewart Cohen, "Knowledge, context, and social standards," *Synthese* 73 (1987), 3–26
- Stewart Cohen, "How to be a fallibilist" in James Tomberlin, ed., *Philosophical Perspectives* 2 (1988), 91–123
- Stewart Cohen, "Two kinds of skeptical argument," *Philosophy and Phenomenological Research* 58 (1998), 143–59
- Stewart Cohen, "Contextualism, skepticism, and the structure of reasons" in James Tomberlin, ed., *Philosophical Perspectives* 13 (1999), 57–90
- Jonathan Dancy, *Introduction to Contemporary Epistemology* (Blackwell, Oxford, 1985)
- Martin Davies, "Externalism, architecturalism, and epistemic warrant" in Crispin Wright, Michael Smith and Cynthia Macdonald, eds., *Knowing Our Own Minds: Essays in Self-Knowledge* (Oxford Univ. Press, Oxford, 1998), 321–61
- Martin Davies, "Externalism and armchair knowledge" in Paul Boghossian and Christopher Peacocke, eds., *New Essays on the A Priori* (Oxford Univ. Press, Oxford, forthcoming)
- Daniel Dennett, "Quining qualia" in Anthony Marcel and E. Bisiach, eds., *Consciousness in Contemporary Science* (Oxford Univ. Press, Oxford, 1988)
- Daniel Dennett, *Consciousness Explained* (Little, Brown & Co., Boston, 1991)
- Keith DeRose, "Solving the skeptical problem," *Philosophical Review* 104 (1995), 1–52
- Fred Dretske, *Naturalizing the Mind* (MIT Press, Cambridge, Mass., 1995)
- Gareth Evans, *The Varieties of Reference*, ed. John McDowell (Oxford Univ. Press, Oxford, 1982)
- Richard Feldman, "In defense of closure," *Philosophical Quarterly* 45 (1995), 487–94
- Richard Feldman and Earl Conee, "Evidentialism," *Philosophical Studies* 48 (1985), 15–34
- Alvin Goldman, *Epistemology and Cognition* (Harvard Univ. Press, Cambridge, Mass., 1986)
- Gilbert Harman, *Thought* (Princeton Univ. Press, Princeton, 1973)
- Gilbert Harman, "The intrinsic quality of experience" in James Tomberlin, ed., *Philosophical Perspectives* 4 (1990), 31–52
- Peter Klein, *Certainty: a refutation of skepticism* (Univ. of Minnesota Press, Minneapolis, 1981)

- Peter Klein, "Skepticism and closure: why the evil genius argument fails," *Philosophical Topics* 23 (1995), 213–36
- Saul Kripke, "Naming and necessity" in Donald Davidson and Gilbert Harman, eds., *Semantics of Natural Language* (Reidel, Dordrecht, 1972), 252–355
- William Lycan, *Consciousness and Experience* (MIT Press, Cambridge, Mass., 1996)
- Colin McGinn, *Mental Content* (Blackwell, Oxford, 1989)
- G.E. Moore, "Hume's philosophy," *The New Quarterly* (1909). Reprinted in Moore (1922), 147–67
- G.E. Moore, *Philosophical Studies* (Routledge and Kegan Paul, London, 1922)
- G.E. Moore, "Proof of an external world," *Proceedings of the British Academy* 25 (1939), 273–300. Reprinted in Baldwin, ed. (1993), 147–70
- G.E. Moore, "Certainty," in *Philosophical Papers* (George Allen & Unwin, London, 1959). Reprinted in Baldwin, ed. (1993), 171–96
- Robert Nozick, *Philosophical Explanations* (Harvard Univ. Press, Cambridge, Mass., 1981)
- Christopher Peacocke, *Sense and Content* (Oxford Univ. Press, Oxford, 1983)
- Christopher Peacocke, *A Study of Concepts* (MIT Press, Cambridge, Mass., 1992)
- George Pitcher, *A Theory of Perception* (Princeton Univ. Press, Princeton, 1971)
- Alvin Plantinga, *Warrant: The Current Debate* (Oxford Univ. Press, New York, 1993a)
- Alvin Plantinga, *Warrant and Proper Function* (Oxford Univ. Press, New York, 1993b)
- John Pollock, *Knowledge and Justification* (Princeton Univ. Press, Princeton, 1974)
- John Pollock, "Reliability and justified belief," *Canadian Journal of Philosophy* 14 (1984)
- John Pollock and Joseph Cruz, *Contemporary Theories of Knowledge*, 2nd ed. (Rowman & Littlefield, Totowa, NJ, 1999)
- James Pryor, "Highlights of Recent Epistemology," forthcoming in *British Journal for the Philosophy of Science*
- John Searle, *Intentionality* (Cambridge Univ. Press, Cambridge, 1983)
- Sydney Shoemaker, "Qualities and qualia: what's in the mind?," *Philosophy and Phenomenological Research* Suppl. Vol. 50 (1990), 109–31
- Sydney Shoemaker, "Self-knowledge and 'inner sense'," *Philosophy and Phenomenological Research* 54 (1994), 249–314
- Roy Sorensen, "Dogmatism, junk knowledge, and conditionals," *Philosophical Quarterly* 38 (1988), 433–54
- Ernest Sosa, "Beyond skepticism, to the best of our knowledge," *Mind* 97 (1988), 153–88
- Barry Stroud, *The Significance of Philosophical Scepticism* (Oxford Univ. Press, Oxford, 1984)
- Michael Tye, "Visual qualia and visual content" in Tim Crane, ed., *The Contents of Experience* (Cambridge Univ. Press, Cambridge, 1992), 158–76
- Michael Tye, "Qualia, content and the inverted spectrum," *Noûs* 28 (1994), 159–83
- Michael Tye, *Ten Problems of Consciousness* (MIT Press, Cambridge, Mass., 1995)
- Thomas Tymoczko and Jonathan Vogel, "The exorcist's nightmare: a reply to Crispin Wright," *Mind* 101 (1992), 543–52
- Jonathan Vogel, "Are there counter-examples to the closure principle?" in Michael D. Roth and Glenn Ross, eds., *Doubling: Contemporary Perspectives on Skepticism* (Kluwer, Dordrecht, 1990a), 13–27
- Jonathan Vogel, "Cartesian skepticism and inference to the best explanation," *Journal of Philosophy* 87 (1990b), 658–66
- Crispin Wright, "Strict finitism," *Synthese* 51 (1982). Reprinted in Wright (1993), 107–75
- Crispin Wright, "Facts and certainty," *Proceedings of the British Academy* 71 (1985), 429–72
- Crispin Wright, "Rule-following, meaning and constructivism," in Charles Travis, ed., *Meaning and Interpretation* (Blackwell, Oxford, 1986), 271–97
- Crispin Wright, "Skepticism and dreaming: imploding the demon," *Mind* 100 (1991), 87–115
- Crispin Wright, *Realism, Meaning and Truth*, 2nd ed. (Blackwell, Oxford, 1993)
- Crispin Wright, "Cogency and question-begging: some reflections on McKinsey's paradox and Putnam's Proof," *Philosophical Issues (Yearly Supplement to Noûs)* 10 (2000), 140–163