

WHAT'S WRONG WITH MOORE'S ARGUMENT?

James Pryor*
Princeton University

1 Some Diagnoses

Moore looked at his hands and argued:

- (1) Here are two hands.
- (2) If hands exist, then there is an external world.
- (3) So there is an external world.¹

Something about this argument sounds funny. As we'll see, though, it takes some care to identify exactly what Moore has done wrong.

I will assume that Moore knows premise (2) to be true. One could inquire into *how* he knows it, and whether that knowledge can be defeated; but I won't. I'll focus instead on what epistemic relations Moore has to premise (1) and to his conclusion (3).

It may matter which epistemic relations we choose to consider. Some philosophers will diagnose Moore's argument using Contextualist machinery. They'll say:

In some contexts, it'd be true to count Moore as knowing he has hands and that there is an external world. In more restrictive contexts, it would not. Moore's argument sounds funny because Moore plays fast and loose with the context. His straightforward assertion "Here are two hands" invites us to occupy a lax context; but one would only be *concerned to argue* that there's an external world in a more restrictive context, where the existence of the external world is an open question.

If these philosophers are right to count 'knows' as context-sensitive,² then they're probably right in their complaint that Moore's performance plays

fast and loose with the context. However, I don't think that can be a full account of the epistemology of Moore's argument. Why? Because some epistemological predicates resist Contextualist treatment. Even if you think 'knows' is context-sensitive, you'll probably grant that predicates like:

... gives Moore some justification to believe ...
 ... makes Moore more justified in believing ... than he was

are not. And the funny epistemic qualities of Moore's argument are manifest even when we confine ourselves to those predicates. Pre-reflectively, it seems like Moore's perception of his hands should give him more justification to believe he has hands than he'd have without it. And we grant that hands are external objects (and that Moore knows them to be so). Yet many are reluctant to accept that Moore's perception of hands gives him more justification to believe there's an external world. Why do we hesitate? If something gives you justification to believe P, and you know P to entail Q, then shouldn't it give you justification to believe Q, too?

Some philosophers will say that Moore's argument just illustrates that Closure is false: you can have justification to believe P, know that P entails Q, and yet fail to have justification to believe Q. But recent years have seen vigorous defenses of Closure; and my sympathies lie largely with the defenders.³ Minimally, I'll assume that if Moore has some justification to believe he has hands, he also has (at least as much) justification to believe there's an external world. We need to say more about this, but for the moment let's press on.

Three diagnoses of why Moore's argument sounds so unconvincing remain.

One diagnosis is the skeptic's. He'll agree that *if* Moore's perceptual experiences gave him justification to believe he has hands, they'd also give him justification to believe there's an external world. But he denies that Moore's perceptual experiences give him any justification to believe *either* hypothesis. That's why Moore's argument generates no conviction.

A second diagnosis has been advanced by Martin Davies and Crispin Wright. They say that Moore *does* acquire perceptual justification to believe he has hands, but this justification doesn't "transmit" across the entailment from (1) to (3).⁴ Davies' and Wright's point is not that Moore *lacks* justification to believe (3). They'll allow that the relation of *having some* justification is closed across known entailment. But they'll point out that *not every* epistemic relation is so closed. (For instance: you might *know* P *non-inferentially*, and know that P entails Q, without knowing Q *non-inferentially*.) So it's a good question whether the relation of *giving you more* justification is closed across known entailment.

As it turns out, it's not. Suppose you start with its being 80% likely for you that Clio's pet is a dog. Then you're informed that Clio's pet has no hair. One effect of this information is to raise the likelihood that her pet is an

American Hairless Terrier, which hypothesis *entails* that it's a dog. But the information also *decreases* the total likelihood that Clio's pet is a dog. It makes it more likely that she owns a fish or a bird. So: evidence can give you more justification to believe P than you had before, you can know P to entail Q, and yet your evidence make you *less* justified in believing Q than you were before.

That being noted, I doubt that this example provides a useful analogy for thinking about Moore. Moore's situation seems rather different from our situation regarding Clio's pet. (Moore doesn't learn that he has hands in an unexpected way that makes it *less* likely that there's an external world.) Davies and Wright will argue, though, that there's *something* about the structure of Moore's justification for (1) and (3) that prevents his evidence for the former from adding credibility to the latter. They'll argue that Moore's experiences of hands are *only able to* justify him in believing (1) to the extent that *he has antecedent reason* to believe (3). And they'll say this is why Moore's argument sounds so unconvincing. We'll examine this diagnosis closely in sections 3 and 4.

I will be defending a third diagnosis. I think that Moore *does* have perceptual justification to believe he has hands, and I think his justification to believe that *does* transmit to the hypothesis that there's an external world. So Moore *can* acquire justification to believe there's an external world by having experiences of hands and reasoning in the way he does. The challenge is to explain our squeamishness about his argument in some way that respects these claims.

2 Theory of Justification

It's essential for insight to these issues to get clear about what we're evaluating. Are we evaluating *the proof* Moore rehearsed? *The reasoning* he engaged in? or what?

In all, I discern five targets of evaluation. They are importantly different. I'll identify two of them now, and introduce the others at later stages in our discussion.

First, there is the **proof** Moore rehearsed: a sequence of propositions and derivation-rules. There doesn't seem to be anything objectionable about this. The proof is clearly valid, and all but the skeptic will grant that its premises are true.

A second, and more interesting target of evaluation is the **justificatory structure** that Moore seems to be endorsing. Does Moore's argument articulate a structure your justification genuinely can have? Can the credibility of (3) really be *enhanced by* your perceptual justification to believe (1)?

Wright, Davies, and the skeptic agree that justification *cannot* be structured in that way. I'll be arguing that it can. I think you genuinely *do* get justification to believe the external world exists from your perceptual justification to believe hands exist. So if what we're evaluating is the justificatory structure that Moore proposed, I think that there is nothing objectionable here, either. Defending this will be my first task.

I'll begin with some groundwork in the theory of justification. I understand **justification** to be the quality that hypotheses possess for you when they're epistemically likely for you to be true, and so epistemically appropriate for you to believe. That's not meant as an *explanation*: if you're not already familiar with the quality I'm talking about, then you won't understand talk about epistemic likelihood, either. Rather, I'm trying to *specify which of various uses* of 'justification' I'm employing. Some philosophers use "X has justification to believe P" or "X is justified in believing P" to mean merely that X is epistemically blameless for believing P.⁵ Some use it to mean that X has a reflective appreciation of *why* P is appropriate for him to believe, and so can offer arguments for P. Given the way I'm using 'justification,' those are all substantive claims; and as it happens, I think they're false. I think you can be careful and blameless in managing your beliefs, and still believe things that aren't well-supported by your evidence—and so aren't epistemically likely for you to be true. The mistakes you're making may be too subtle and well-entrenched for you to recognize. So being blameless for believing P doesn't guarantee genuinely having *justification* to believe P. Neither does having justification to believe P require you to *have some argument* you could present for P. Justification need not be that sophisticated and reflective.⁶

I said that justification is a quality that *hypotheses* can possess for you: they can be hypotheses you have *justification to believe*. It doesn't matter whether *you do* believe them. This is sometimes called "propositional justification." We also have a notion of *doxastic attitudes* being justified. I will return to that notion in section 5. For now, let's keep thinking about propositional justification.

Justification comes in degrees and it can be defeated. It can be defeated in different ways. Suppose my brother tells you that his landlord is shifty-looking. That gives you some justification to believe that his landlord is dishonest. One way for that justification to be defeated is for my brother's roommate to tell you that their landlord is *not* shifty-looking. That evidence *opposes* my brother's testimony. It gives you some justification to believe the opposite. Another way for your justification to be defeated is for you to learn that my brother's landlord is an active member of his church and donates generously to charity. This evidence *narrows the reference class*. Shifty-looking people are in general likely to be dishonest, but shifty-looking people who are active in their church and so on tend not to be. A third way for your justification to be defeated is for me to tell you that my

brother never met his landlord, and just has a prejudice against him because of a disagreement over the rent. This evidence attests to my brother's *not being in a position to know* what his landlord looks like. So it *undermines* the justification my brother's testimony gave you to believe that the landlord is dishonest. It doesn't give you any special reason to believe the landlord *is* or *looks honest*; he *may very well* look shifty and be dishonest. My testimony just gives you less reason to rely on my brother's word for it.

Of course, since justification comes in degrees, all these varieties of defeating evidence will come in degrees too. And they may themselves be defeated or undermined by further considerations.

Since justification to believe P can be defeated in these ways, we need to distinguish prima facie justification from all things considered justification. By "prima facie justification" to believe P, I *don't* mean merely that at first glance, *it seems like* you have justification to believe P. You really must have a body of justification for P. That justification can at the same time *be defeated or undermined* by further evidence you possess; but *in the absence of* such further evidence—and when the further evidence *itself* gets defeated—your prima facie justification to believe P must constitute some degree of *all things considered* justification to believe P.⁷ That's what I mean by "prima facie justification" to believe P.

There are several different roles a theory of justification can assign a hypothesis H. To begin with, the theory needs to identify conditions whose *truth* is what *makes you* have prima facie justification to believe P. Let M be those conditions: conditions the mere satisfaction of which is supposed to make you justified. (You don't need to be *aware* that you satisfy them—unless that too is one of the conditions included in M.) One role a theory of justification can assign a hypothesis H is to include *H's being true* in M. That is, in order for you to have a given kind of prima facie justification to believe P, H has to be true. I'll call this a **truth-requiring** treatment of H.

Various hypotheses U will be such that evidence for them *undermines* the prima facie justification that satisfying M gives you to believe P. Hence, in order to be *all things considered* justified in believing P, you'll need to *lack* (undefeated) justification to believe that U obtains.⁸ In the example I gave several paragraphs ago, what made you prima facie justified in believing the landlord is dishonest included hearing my brother testify that he's shifty-looking, and knowing that shifty-looking people are likely to be dishonest. One undermining hypothesis U for that justification is that my brother has never seen his landlord and so doesn't know what he looks like. My testimony gave you some justification to believe U, and thus (to some degree) undermined the prima facie justification you got from listening to my brother.

We're now ready to distinguish two further roles a theory of justification can assign a hypothesis H. A theory treats H **conservatively** when it says that you need some justification to believe H *in order to* have a given kind of

prima facie justification to believe P. That is, the conditions M that *make you* have that prima facie justification include your having this justification to believe H. The justification to believe H has to come from sources *other than* the justification to believe P that we're considering, since it needs to be in place as a precondition of *your having* that justification to believe P. I'll put this by saying that your justification to believe H needs to be **antecedent** to this justification to believe P.⁹ (It's *allowed* to derive from justification to believe P you have by other routes.)

A theory that treats H **liberally** *denies* that having prima facie justification to believe P requires you to have antecedent justification to believe H. But it does count not-H as an *undermining* hypothesis: evidence against H undermines your prima facie justification to believe P.¹⁰

Conservative and liberal treatments of H may or may not also assign H a truth-requiring role. If a theory assigns H a truth-requiring role, but *neither* of the other roles, I'll call that a *merely* truth-requiring treatment of H.¹¹

A typical theory of justification will choose different options for different hypotheses. It might be conservative about some hypotheses, liberal about others, and merely truth-requiring about still others.

Let's consider some examples.

Suppose you're reading some proof of the Pythagorean Theorem. H1 is the claim that you *understand* and *correctly follow* the proof. Presumably, for you to be justified in believing the theorem, H1 does have to be true. But you don't need to have *evidence that* H1 is true. It's the proof itself that justifies you in believing the theorem. H1 is just some condition that *enables* this to happen.¹² It's not itself one of the premises that your justification for believing the theorem rests on—not even a suppressed, background premise. So the right treatment of H1 seems to be a non-conservative one.¹³

A second example. You remember parking your car in Lot 15. After you finish teaching, you intend to walk to Lot 15 and drive your car home. Presumably, having this intention gives you some justification to believe *you will* walk to Lot 15 and drive your car home. However, that assumes that your car is *still located* in Lot 15, and hasn't been stolen or towed away. Let H2 be the hypothesis that your car is still there. In this case, a conservative treatment seems most plausible. Having the intention to go drive your car home is not enough *by itself* to give you prima facie justification to believe you'll succeed in doing it. You also need some antecedent justification to believe H2 is true—that your car will still be there when you arrive.

A third example. You have visual experiences of your car. H3 is the claim that those experiences are reliable. A *reliabilist* will take H3 to have a status like H1's: it has to *be true*, for your experiences to justify you in believing your car is present, but you don't need to have *evidence or justification to believe* it's true. Internalists, on the other hand, deny that H3 needs to be true. Either they'll treat H3 liberally, and say it's enough that you *lack* reason to believe your experiences are *unreliable*. Or they'll treat it

conservatively, and say you do require antecedent reason to believe your experiences are reliable. But in neither case does *the truth* of H3 make an epistemological difference. It's only *your epistemic situation concerning H3* that is important.

As we can see, one will likely handle different cases differently.

3 Perceptual Justification

It seems to Moore that he's perceiving his hands. There are various things he could learn that would entail he's not. I'll call these **non-perceiving hypotheses**: they're hypotheses that are (known to be) incompatible with his experiences being genuine perceptions. For instance, the hypothesis that Moore is hallucinating his hands is a non-perceiving hypothesis. It's compatible with Moore's *having* hands; but it denies that he's perceiving them. Evidence for non-perceiving hypotheses will tend to *undermine* any justification Moore's experiences give him to believe he has hands.¹⁴

What role should we assign these hypotheses? Should we treat them conservatively, and say that Moore needs to have antecedent justification against them, as part of the conditions that *make him* have any prima facie perceptual justification? Or should we treat the hypotheses liberally, and say they *just count as underminers*?

A **conservative about perception** treats all non-perceiving hypotheses conservatively.¹⁵ In my 2000, I defend a view that treats all such hypotheses liberally. I call that view **dogmatism about perception**. Intermediate views are also possible, but I will focus on these two.

I understand Wright to be a conservative about perception. This attribution is a bit complicated, though. One difficulty is that he uses different terminology than I use. The epistemic quality I call "justification" includes what he calls "entitlements" and "warrants."¹⁶ A second difficulty is that I'm presently discussing propositional justification—justification *to believe* certain hypotheses—whereas Wright is often concerned whether you really have *justified beliefs*, and if so, what processes were involved in your acquiring them. I want to reserve questions about doxastic justification for later. Doxastic justification is complicated. It will pay to get clear about propositional justification first. A third difficulty is that Wright sometimes shifts to higher-order questions: not what it would take for you to justifiably believe *P*, but what it would take for you to justifiably believe that *you have justification* for *P*, or that *you're perceiving* that *P*. These higher-order questions are difficult, too. Some of our discussion will bear on them (see note 33 and section 7). But mostly, I'll stay focused on questions about your first-order justification to believe *P*, and I'll interpret Wright's views as directed to those questions, too.¹⁷ A final difficulty is that Wright's position has evolved recently, in ways that complicate the dialectic

between us (see his 2004). To keep our discussion manageable, I'll ignore the latest developments.

If we translate Wright's pre-2004 views into our present framework, this is what he says: In order for Moore's experiences to give him any justification to believe he has hands, Moore does need to have antecedent justification to believe that he's not hallucinating, that he's not a brain in a vat, that there's an external world, and so on.¹⁸ As it happens, Wright thinks that Moore *has* that justification. It's not justification that Moore did anything special to earn. It didn't require him to engage in any a posteriori (or a priori) inquiry. He gets it by default. We all have defeasible default justifications to believe that we're not hallucinating, that there's an external world, and so on.¹⁹ Our perceptual beliefs about hands aren't typically *inferred from* or based on prior beliefs that we're not hallucinating; and they don't *need* to be. We may just tacitly *assume* we're not hallucinating. But Wright thinks we *do* need antecedent *justification* for that assumption, before our experiences will even *prima facie* justify our perceptual beliefs.²⁰

My view is that when Moore's experiences represent there to be hands, that *by itself* makes him *prima facie* justified in believing there are hands. This justification doesn't rest on any *premises about* Moore's experiences: whether they constitute perceptions, how reliable they are, or anything like that. It's in place so long as he *merely has* experiences that represent there to be hands. There are things Moore could learn that would undermine this justification. But it's not a condition for having it that he *first* have justification to believe those undermining hypotheses are false.

That is the view I defended in Pryor 2000. I emphasized there how simple and intuitively appealing the view is; I said the main work for systematic epistemology should be to defend the view against challenges. I think I was overly modest. But systematic epistemology can do more. One way it can do more is by highlighting how *unconvincing* other answers to skepticism are. In fact, I suspect we'll have *no* prospect of avoiding skepticism unless we agree to be liberal to *some* degree, about *some* kinds of hypotheses. This puts pressure on any non-skeptic to say what's wrong with *the particular kind* of liberalism that the perceptual dogmatist espouses.

Another way for systematic epistemology to do more is to give a *positive* account of why our perceptual experiences should have the epistemic powers the dogmatist says they have. What I've said so far is compatible with a variety of such accounts. Some will argue that experiences as of P justify you in believing P because they make that belief *irresistible*, and it can't be the case that you ought not believe what you can't help believing.²¹ Others will argue there's something distinctive about *the concepts* we employ in our perceptual beliefs, which makes those beliefs epistemically appropriate responses to our experiences.²² I'm not sympathetic to either of those approaches. My view is that our perceptual experiences have the epistemic powers the dogmatist says they have because of what the *phenomenology* of perception is like.

I think there's a distinctive phenomenology: the feeling of *seeming to ascertain* that a given proposition is true. This is present when the way a mental episode represents its content makes it feel as though, *by enjoying* that episode, you can *thereby just tell* that that content obtains. We find this phenomenology in perception and in memory. When you have a perceptual experience of your hands, that experience makes it feel as though you *can just see* that hands are present. It feels as though hands are being shown or revealed to you. This phenomenology may be present in other mental episodes, too. But it's not present in *every* representational mental episode. When you daydream or exercise your visual imagination, you represent propositions (the same propositions you represent when you perceive), but it does not feel as though you can thereby just tell that those propositions are true.²³

My view is that our perceptual justification comes from that phenomenology. Having the phenomenology of seeming to ascertain P is *what makes us* have prima facie justification to believe P. (You really need *to have* the phenomenology. It's not enough to think you do.)

Often, our perceptual experiences will move us to believe *more* than just what is presented in the phenomenologically distinctive way I described. For example, you may enjoy the phenomenology of seeming to ascertain that a blue-uniformed man is present, and unreflectively believe in response that the police have arrived. As I'm imagining this case, you *don't* have the phenomenology of seeming to ascertain that the police have arrived.²⁴ In such a case, I don't think your experiences are enough, *by themselves*, to give you prima facie justification to believe the police have arrived. They only justify you in believing the propositions they give you the phenomenology of seeming to ascertain.²⁵ To get justification to believe that the police have arrived, you'd need further justification to believe that blue-uniformed men are likely to be members of the police.

It's not easy to discern what propositions we "seem to ascertain," and what propositions we merely unreflectively infer. Our *perceptual reports* don't track the difference very closely. When you have the right kinds of background evidence, you'll unhesitatingly *say* things like "It looks as if the police have arrived" and "I see that the Smiths have already left for Australia." But I think there is a real difference. And my theory is that what we seem to ascertain, we thereby have immediate prima facie justification to believe. We don't need to have antecedent justification to believe we're not hallucinating, and so on.

4 Some Different Types of Epistemic Dependence

I've described two epistemologies of perception: the conservative and the dogmatist. They disagree about Moore's justification to believe (1).

According to the conservative, Moore's experiences give him justification to believe (1) only if he has antecedent justification to believe (3). According to the dogmatist, Moore's experiences immediately justify him in believing (1) and don't require him to have antecedent justification to believe anything else.

In particular cases, a conservative like Wright and I will largely agree about *which* subjects have justification. Wright says subjects need antecedent justification against undermining hypotheses, but they get it by default. I say they don't need it. We'll disagree about some details. For example, we'll disagree about when subjects have justification to believe the undermining hypotheses are false. But our central disagreement isn't about *who has* justification; it's about *the structure* of their justification. (Compare: the central disagreement between modal realists and ersatzists is not about *which* propositions are possible; it's about what their possibility consists in.) As we'll see now, this disagreement about justificatory structure drives further disagreement about which arguments count as objectionably question-begging.

Moore's justification for (1) is often said to "presuppose" or already "epistemically depend" in some way on his conclusion (3). This is thought to prevent his justification for (1) from lending any additional credibility to (3). He can't *get* any justification to believe (3) from his perceptual justification to believe (1). In Wright's and Davies' terminology, Moore's justification for (1) doesn't **transmit to** his conclusion (3).

Let's figure out what kind of "epistemic dependence" Moore's argument really exhibits, and whether that dependence really does have the claimed effect. I'll canvas five types of dependence; only the fourth and fifth will hold our interest.

Type 1. One way for an argument's premise to epistemically depend on its conclusion is that, in order for the premise to be true, it's necessary that the conclusion is true too. But there can't be anything epistemologically objectionable about *that*; if there were then deductive arguments would never be legitimate.

Type 2. Another way for a premise to epistemically depend on a conclusion is that, in order for you to *have justification to believe* the premise, it's necessary that the conclusion is true.²⁶ One problem here is that, again, it's not clear there's anything epistemologically objectionable about it. Suppose you're wondering whether people have justification to believe *anything*. You persuade yourself that if nothing else, at least cogito judgments are justified. That is, you reason:

- (4) Whoever believes he exists is justified in so believing, since his belief must be true.
- (5) I believe I exist.
- (6) So I'm justified in believing I exist.

- (7) So there are at least some cases of people having justification.

This argument exhibits the type of epistemic dependence we're considering. It's necessary for you to have justification to believe *any* of its premises that its conclusion be true. But—setting aside controversies about why premise (4) should be true—this argument seems perfectly respectable. It seems like a perfectly good way to argue that at least some beliefs are justified.²⁷

Furthermore, it's not clear that this type of epistemic dependence is one that Moore's argument even *exhibits*. Many epistemologists would allow Moore to be justified in believing he has hands even if *he were* an immaterial spirit, and the external world a mere hallucination. So he could have justification for his premises, even if his conclusion were false.

Let's keep looking, then.

Type 3. What about arguments where for you to have justification to believe the premise, it's necessary that *you have justification to believe* the conclusion? Once more, this seems to include some arguments that are perfectly respectable. It includes arguments where the connection between premise and conclusion is *so obvious* that understanding the premise well enough to be justified in believing it *requires* you to take any justification for the premise to also justify you in believing the conclusion.

A difficulty common to the proposals we've considered so far is that they characterize the premise's dependence on the conclusion as some kind of **necessary condition**. That doesn't seem to be what we need. Let's try a different approach.

Type 4. Another type of dependence between premise and conclusion is that the conclusion be such that evidence *against it* would (to at least some degree) undermine the kind of justification you purport to have for the premises.²⁸ Moore's argument clearly *does* exhibit this type of dependence. So long as we maintain the assumption that hands are external objects, any evidence that *there is no external world* will (to some degree) undermine Moore's perceptual justification for believing he has hands.

But is this type of dependence, in itself, a bad thing?

That's a difficult question, because many arguments that exhibit it will also exhibit a further type of epistemic dependence.

Type 5. We have this type of dependence when having justification to believe the conclusion *is among the conditions that make you* have the justification you purport to have for the premise. That is, whenever you need *antecedent* justification to believe the conclusion, as condition for having that justification for the premise.

Type 5 dependence does clearly seem to be an epistemic vice. Consider an example we discussed before:

- (8) I intend to walk to Lot 15 and drive home.
(9) So *I will* walk to Lot 15 and drive home.

(10) So my car will still be in Lot 15 when I get there.

This argument sounds bad to us because we think your intention is not enough, by itself, to justify you in believing you'll succeed in driving your car out of the lot. We think you also need antecedent justification to believe your car is still *in* the lot. Wright and Davies argue that any Type 5 argument will **fail to transmit** the kind of justification you have for its premises to its conclusion.²⁹ Your justification for (9) *relies on* antecedent justification to believe (10); and for that reason, it can't *make* (10) any more credible for you.

I agree that Type 5 dependence ruins an argument. But what about Type 4 dependence?

There are several questions to address:

- Q1. Is it *possible* for an argument to exhibit Type 4 dependence while failing to exhibit Type 5?
- Q2. If so, are arguments that *merely* exhibit Type 4 dependence epistemologically objectionable?
- Q3. In which group does Moore's argument fall?

Let's take Q1 first. If you insist on treating *all* undermining hypotheses conservatively, then you can maintain that the two types of dependence coincide. Whenever evidence against a conclusion would undermine your justification to believe its premise, you'll think your justification for the premise must already rest on, and require antecedent justification for, the assumption that the conclusion is true.

But if you're willing to be liberal about *any* undermining hypotheses, then you think there can be undermining hypotheses that you don't need to be antecedently able to eliminate. This opens up room for arguments that have Type 4 dependence but not Type 5.

Which arguments *exhibit* that dependence will depend on what hypotheses you're willing to be liberal about. But here's a plausible example. Your introspective beliefs about what sensations you're having are fallible. You can be primed to expect sensations of cold and actually be given sensations of heat. In such cases you'll believe that you're having sensations you're not having.³⁰ So the hypothesis that you're making a priming mistake looks like an underminer for your introspective justification for believing you feel cold. Evidence that *you are* making a priming mistake ought to diminish the credibility of your introspective belief by at least *some degree*. At the same time, it's not plausible that your justification to believe you're having a given sensation requires you to have antecedent justification to believe you're not making any priming mistakes. *Sophisticated* subjects may know that they're reliable about their sensations. But I think you can have justified beliefs about your sensations long before attaining that degree of

epistemic sophistication. So the hypothesis that you're making a priming mistake is not one you need antecedent justification to rule out.

Suppose that's all correct. Now consider a case where you *genuinely have* a cold sensation, are aware of having it, and you believe you have it. On the basis of your introspective awareness of your sensation, you judge that you're really having the sensation you think you're having, so you're not making a priming mistake right now.³¹

That piece of reasoning seems to exhibit Type 4 dependence, without exhibiting Type 5. It also sounds to me like an *epistemologically respectable* piece of reasoning. I think your introspective awareness of your sensation does make the hypothesis that you're not making a priming mistake *somewhat* more credible. You needn't be a dogmatist about perception to agree. You only need to be willing to treat this one undermining hypothesis for introspective justification in a liberal way.

In answer to Q1, then, I think yes, there can be arguments that exhibit Type 4 dependence without exhibiting Type 5. I think any liberal should be open to this possibility. I also think such arguments can be epistemologically respectable.³² I hope the reasoning I just described gives a useful example; but I'll need to do a lot more to make the claim fully plausible. The next three sections will try.

Since I'm a dogmatist about perception, I think Moore's argument is another case where we have Type 4 dependence without Type 5. Here are a few more. Suppose you're watching a cat stalk a mouse. Your visual experiences justify you in believing:

(11) The cat sees the mouse.

You reason:

(12) If the cat sees the mouse, then there are some cases of seeing.

(13) So there are some cases of seeing.

Evidence against (13) would undermine your visual justification to believe (11); but I don't think you need antecedent justification to believe (13), before your experiences can give you justification to believe (11). I *also* think it's plausible that your perceptual justification to believe (11) contributes to the credibility of (13).

Suppose you look at a wall that's been painted red. Your visual experiences justify you in believing:

(14) The wall is red.

You reason:

- (15) If the wall is red, it's not white but lit by tricky red lights that make it appear red.
- (16) So the wall is not white but lit by tricky red lights.

Here too I think your visual justification to believe the premise makes the conclusion more credible for you. Your justification to believe the wall is red contributes to the credibility of the claim that the wall isn't white but lit by tricky red lights.³³

I think all of these arguments are *epistemologically* respectable: that is, they articulate structures your justification genuinely can have.

When people learn this is my view, they complain that I'm giving away too cheaply justification to believe we're not in undermining scenarios. Stewart Cohen has argued this forcefully in Cohen 2002. One of his arguments involves a dialogue with his son. His son asks whether a certain table is red, and Cohen replies, "Yes I can see that it's red." The son asks whether Cohen knows that it's not white but lit by tricky red lighting, and Cohen imagines replying with an argument like my Red Wall argument (14)–(16). Cohen complains: "Surely [my son] should not be satisfied with this response."³⁴

I agree that there are *some* respects in which these arguments are persuasively crippled, and so can fail to satisfy. I'll try to characterize those respects in the remaining sections of this paper. In terms of their justificatory structure, though, I think these arguments have nothing to be ashamed of.

5 Reasoning and Doubt

I've argued that there's nothing wrong with the justificatory structure that Moore's argument articulates. It's a structure that your justification genuinely can have.

Let's now take up a new target of evaluation: the **reasoning** that Moore engaged in.

We have justification to believe many things that we don't yet believe. Reasoning is a process by which we actively try to bring ourselves around. We try to believe what we have justification to believe. Usually this will involve recognizing that some beliefs you already have make a conclusion likely to be true, and believing the conclusion in response.

What does it mean to believe a conclusion "in response"? That depends. Sometimes reasoning just aims to "tease out" implications of things you already believe,³⁵ and undertakes no commitment about the epistemic priority of premises and conclusion. For example, suppose you see Joey emerge from a classroom with a big smile on his face. You believe that he's performed well on his Latin exam. *Your justification* to believe this will come

in part from your having justification to believe that *he had* a Latin exam today. But *your first explicit recognition* that Joey had a Latin exam today may come by inferring it from your belief that he just performed well on it. You may be thinking: "Look how happy Joey is. He must have aced the Latin exam I see him emerging from. Wait a second! If Joey *just aced* the Latin exam, that means he *just took* the Latin exam! Damn it, I thought that exam was on Thursday . . .!"

We sometimes do reason in that "teasing-out" way. But I think often our reasoning aims to do more: it aims to reconstruct *the structure* of our justification. It aims to make explicit the justificatory relationships *by virtue of which* we have justification to believe a conclusion. When you reason from P to Q in *this way*, your resulting belief in Q will be (at least in part) *based on* your belief in P, and your recognition that P supports Q.³⁶ For the rest of our discussion, I'll understand "reasoning" to mean reasoning that has this aim.

I've argued that Moore's *justification can have the structure* he presents it as having. Let's now consider whether *the piece of reasoning* that Moore engaged in can be epistemologically legitimate, too.

One complaint often charged against Moore's argument—and my Red Wall argument—is that anyone who had doubts about its conclusion couldn't *use* the argument to rationally overcome those doubts. I think this complaint is correct. Let's see if we can respect it, while still maintaining that your justification has the structure the dogmatist says it has.

We need to get clear about what "doubting" an argument's conclusion involves. The dogmatist will agree that *having evidence against* the conclusion of Moore's argument undermines your perceptual justification to believe the premise. To get a disputed case, we need to consider examples where you lack such evidence. One possibility is that you start out believing—or at least suspecting—the conclusion is false, without having justification for doing so. (There are other ways to understand "doubt," too. We'll consider them later.)

If you *recognize* that your belief or suspicion about the argument's conclusion is unjustified, but can't help having it, I'll say you have a "pathological doubt." Otherwise your doubt is non-pathological. Let's set the pathological doubts aside for now. I assume you can have doubts that are unjustified without being pathological. Can such doubts affect what you're justified in believing? For example, if you *believe* without evidence that there is no external world, does that belief undermine the justification your experiences give you to believe you have hands? If you *believe* without evidence that the wall is lit by tricky red lighting, does that undermine your justification to believe the wall is red?

To answer these questions, we need to keep two epistemological contrasts firmly in mind.

The first contrast is between what you have justification to believe, and what you're **rationally committed** to believe by beliefs you already have. A

rational commitment is a hypothetical relation between your beliefs; it doesn't "detach." That is, you can have a belief in P, that belief can rationally commit you to believe Q, and yet you be under no categorical requirement to believe Q. Suppose you believe Johnny can fly. This belief rationally commits you to the belief that someone can fly. If you're not *justified* in believing that Johnny can fly, though, you need not have any *justification* for the further belief. You may even have plenty of evidence and be fully justified in believing that *no one* can fly. But your belief that Johnny can fly still rationally commits you to the belief that someone can fly. Given your belief about Johnny, if you refrain from believing that someone can fly, you'll thereby exhibit a rational failing.³⁷

I think we can understand rational commitments like this. Take a belief the subject happens to have, e.g., his belief in P. Consider what would be the epistemic effects of his *having (decisive) justification for* that belief. (It will be important what *other* parts of the subject's evidence we also suppose changed; but I won't try to sort that out here.) If one of the effects is that the subject has decisive justification to believe Q, then his belief in P counts as rationally committing him to the belief in Q—regardless of whether he *really does* have any justification to believe P. (Notice: although from the facts that you believe P, and that belief commits you to believe Q, we can't infer that you have any justification to believe Q, we *can* infer that when we add the further premise that you have justification to believe P.)

For our purposes, it will be useful to weaken and generalize this notion. We can do that by considering the effects of having *degrees* of justification to believe P, and considering *a broader variety* of effects. That lets us introduce the following normative relations. We can say that a subject's belief in P **rationaly supports** those beliefs that *justification for* P would lend some credibility to (even if it doesn't decisively commit the subject to them). We can say that a subject's belief in P **rationaly opposes** those beliefs that justification for P would tell against. Most interestingly, we can say that a subject's belief in P **rationaly obstructs** him from believing Q on certain grounds, when justification for P would undermine the justification those grounds give him for Q. For example, the mere belief that your color vision is defective—whether justified or not—would rationally obstruct you from believing the wall is red on the basis of your color experiences. It would rationally pressure you to place less trust in those grounds on that question.

Because the epistemic effects we're considering are non-monotonic and matters of degree, so too will be the relations of rational obstruction, rational support, and rational opposition. To keep our discussion simpler, though, I will mostly suppress the complications that introduces. But I will allow myself to talk about partial doxastic attitudes—attitudes like doubt, suspicion, and disbelief—standing in these various relations.

I will count a belief as **rational** when it's a belief that none of your other beliefs or doubts rationally oppose or rationally obstruct you from believing.

This makes “being rational” a different quality than having justification. A subject can *have some justification* to believe P, but be unable to *rationally* believe P on the basis of that justification, because of some (unjustified) beliefs and doubts he also has. Consider again your belief that your color-vision is defective. Suppose that this belief is unjustified (but you don't realize it). Because you don't have *justification* to doubt your color vision, I don't think the justification you get from your color experiences will be *undermined*. You'll still *have justification* to believe the wall is red. But your actual doubt will *rationally obstruct you* from relying on your color experiences. It'll prevent you from rationally accepting that justification.

I said we need to keep two contrasts firmly in mind. The second contrast is between having *justification to believe* something, and having a *belief* that is justified or well-founded. Unjustified beliefs and doubts may have no undermining effect on what propositions you have justification to believe; but for your beliefs to be well-founded, it's *not enough* that they be beliefs in propositions you have justification to believe. They also have to be *based on* that justification, and they have to be *rational* beliefs. Suppose you believe P, on the basis of what are in fact good reasons for believing P. But you also have doubts that rationally oppose P, or rationally obstruct you from believing P for the reasons you do. Those doubts will render your belief in P *irrational* even if they don't affect your justification *to* believe it. And if your belief in P *is* irrational, then it can't be a justified or well-founded belief. In this way, then, even unjustified doubts can affect what justified beliefs you're able to have.³⁸

Let's suppose you doubt whether the premises of some argument are true. Or you doubt whether the argument really justifies you in believing its conclusion. Or perhaps you believe something that rationally commits you to doubting those things. In all of these cases, it *won't be rational* for you to accept the argument's conclusion. (At least, not on the basis of that argument. You might have other, independent reasons for believing it.) It doesn't matter whether your doubts are justified. Even unjustified doubts will make it irrational for you to accept the conclusion. Those doubts will prevent the argument from rationally persuading you, until you give them up.

But none of that should cast epistemic discredit on *the argument*. For all I've said, the argument could be any standard proof of the Pythagorean Theorem. You may *in fact have* justification to believe its premises, and just have unjustified doubts about them. The proof would then *still give you* justification to believe its conclusion. Your doubts would prevent you from *rationally accepting* that conclusion, until you gave them up. But this is no fault of the proof's. The fault lies with you, for having doubts you have no good reason to have.

The same holds for any argument that exhibits only Type 4 dependence. A subject doesn't need antecedent justification for the conclusion; but if she *acquires reason to doubt* the conclusion, that will undermine her justification for the premises. It follows that if the subject *just happens to* doubt

the conclusion, without having any reason for doing so, her doubt will rationally obstruct her from believing the argument's premises. Hence, her doubt will make it irrational for her to accept the argument's conclusion (at least, on the basis of that argument). So the argument can't "rationally overcome" her doubt. It doesn't give her a piece of reasoning she can *rationally* accept, while starting from a position of having that doubt. To be sure, that's *some* kind of failing. But it's a deficiency in the argument's *persuasive* power, not in its justificatory structure. We just saw that unjustified doubts can make it irrational for you to accept the conclusions of *perfectly respectable* arguments, like standard proofs of the Pythagorean Theorem. These Type 4 arguments are a special case, where your doubts happen to be about the very proposition that is the argument's conclusion. But that makes no epistemological difference. In both cases, you may very well *be justified* in believing the argument's premises, and the argument may very well *give you justification* to believe its conclusion. Given your doubts, it may not be rational for you to accept the premises, or the conclusion. But that won't be the argument's fault. It'll be your fault, for having doubts where no doubt is justified.³⁹

6 Other Varieties of Doubt

Up to this point, I've been understanding "doubt" as some degree of *disbelief*. However, sometimes doubt is a matter of mere *agnosticism* or *suspended judgment*. We might regard the skeptic's hypotheses as equally likely true as false. That's not enough to count as *believing* those hypotheses.

In some cases, I think agnosticism will be compatible with rationally accepting Moore-type reasoning. (If you do accept the reasoning, that will rationally commit you to give your agnosticism up: it's irrational to believe you have hands on the basis of your experiences while *retaining* agnosticism about whether those experiences count as perceptions.) In other cases, I think agnosticism will have the same effects as positive disbelief (albeit to a lesser degree). I'll have to explain and defend this contrast elsewhere.⁴⁰

There is still another way to understand "doubt" about an argument's conclusion. Sometimes we "entertain doubt" about a hypothesis in a sense that's compatible with *continuing to believe* the hypothesis. This is what goes on when we read Descartes' First Meditation. We don't *really* suspend our beliefs about our surroundings; we just *entertain the hypothesis* that those beliefs are false, and think about what epistemic relations we bear to that hypothesis, and what follows. Call this **hypothetical doubt**. (In conversation, Cohen insists this is the only kind of doubt his son has when he asks whether Cohen knows the table isn't white but lit by tricky red lights.)

Are hypothetical doubts enough to undermine or take away our perceptual justification? Some philosophers say they are. For instance, Wright says that:

Once *the hypothesis is seriously entertained* that it is as likely as not, for all I know, that there is no material world as ordinarily conceived, my experience will lose all tendency to corroborate the particular propositions about the material world which I normally take to be certain.⁴¹

I on the other hand want to treat hypothetical doubt along the same lines I sketched in the previous section.

Let me say a little to set my account up. Suppose you do in fact possess some prima facie justification to believe P. Then a trustworthy philosophical authority tells you you *don't* have any justification (even prima facie justification) to believe P. Now you have negative higher-order justification: justification to believe you *don't* have (even prima facie) justification to believe P.

What will the effect of that higher-order justification be on your first-order prima facie justification? Will there be no effect? Will it undermine that first-order justification (at least to some degree)? Will it mean that you no longer satisfy the conditions *making you* have even prima facie justification to believe P? I'm inclined to say the second; but I've encountered sympathy for each of these answers.

Now suppose that, instead of higher-order *justification*, you merely acquire a negative higher-order *belief*. What then will the effect be? I think the most plausible answer here is that the higher-order belief will rationally obstruct you from believing P.⁴²

So, now, consider: You have a visual experience of your hands. That experience gives you some prima facie justification to believe you do have hands. Initially, you're *inclined to believe* you have hands, on the basis of your experiences.

Then along comes a skeptic. He starts presenting various undermining hypotheses U_1, U_2, \dots . He argues that since you have no antecedent justification *against* those hypotheses, your experiences don't really give you justification—even prima facie justification—to believe you have hands. He argues that if *you did* have justification to believe there's an external world, you ought to be able to rationally persuade him that there is, but you can't.

As it happens, the skeptic's arguments are flawed. Before you encountered him, you *did* have justification to believe you have hands and so on, contrary to what he's claiming. But the skeptic is a smooth dialectician. His arguments sound pretty compelling to you. You don't discern their flaws. When a flawed argument sounds compelling to you, we might say *you're justified* in believing its conclusion—at least until further reflection reveals the flaws.⁴³ If we do say that, then you'll be *justified* in believing what the skeptic tells you. You'll

be justified in believing your experiences don't give you any perceptual justification. In my view, that has the result of *undermining* your first-order perceptual justification. Listening to the skeptic will have undermined some prima facie perceptual justification you really have.⁴⁴

Alternatively, perhaps we should *deny* that you're justified in believing the conclusions of flawed but compelling arguments. The skeptic's arguments merely *persuade* you. They don't persuade you that any of U_1 , U_2 , ... are true; but they do persuade you to believe, or at least give some credence to, the negative higher-order claim that your experiences don't give you perceptual justification. (Or perhaps they just persuade you to believe some things that *rationally commit you* to that conclusion.) In that case, I think the skeptic will have succeeded in rationally obstructing you from believing you have hands on the basis of your experiences. So you'll no longer be able to *justifiably believe* you have hands.

Neither of these outcomes means that the skeptic *was right*. The skeptic makes claims about *all subjects*, even subjects who haven't heard his argument. On the story I just told, those claims are false. But subjects who *do* hear the skeptic's arguments, and are partly taken in by them, do really end up with some of the epistemic difficulties the skeptic says we all suffer from.

Skepticism isn't the truth about all of us, then. It's just a disease that some of us catch. The way to cure the disease is to realize that skepticism *isn't* the truth about all of us: the skeptic's arguments are flawed. So our negative higher-order beliefs are false. When we give those negative higher-order beliefs up, then the prima facie justification our perceptual experiences gave us all along will be undefeated and unobstructed.

7 Persuading the Skeptic

Let's revisit the different targets of evaluation I've identified so far.

It was clear from the beginning that the **proof** Moore rehearsed is perfectly respectable.

I argued in sections 3–4 that the **justificatory structure** Moore was endorsing is a genuine one. Perceptual justification to believe you have hands doesn't require antecedent justification to believe there's an external world; and it can make the latter hypothesis more credible for you.

Sections 5–6 gave us the following results:

- It was already agreed that if you *have evidence* that you're in a skeptical scenario, it will (to some degree) undermine your perceptual justification to believe Moore's premise (1).
- If you *believe or suspect* without evidence that you're in a skeptical scenario, that won't *undermine your justification* for Moore's premise (1). But it will (to some extent) *rationally obstruct you* from believing

that premise on the basis of your experiences. So you won't be able to use Moore's argument to rationally overcome your suspicions.⁴⁵

- I think some cases of agnosticism work the same way. Others don't. I'll argue for this elsewhere.
- In the happy case where you neither *have* nor *have reason to have* the kinds of doubts the skeptic wants to induce, then the justification your experiences give you for Moore's premise (1) will be undefeated and unobstructed. Having that justification for the premise *will* make Moore's conclusion more credible for you; and that justificatory relationship is one that you can rationally endorse in your reasoning.

If Moore's psychological and epistemic situation was the last one, then I claim the reasoning he engaged in was perfectly legitimate.

Let's now consider a fourth target of evaluation. Instead of Moore's own reasoning, let's consider the **dialectical power** of his argument.

Arguments are dialectical creatures, in a way that proofs and pieces of reasoning are not. Arguing involves *offering* pieces of reasoning *to audiences*. (You can argue with yourself. I think of that as arguing with a *hypothetical* audience.) An argument *succeeds* for a given audience to the extent that it presents the audience with a piece of reasoning they can rationally accept. I call arguments that do that **dialectically effective**. (I don't care whether the audience *really does* accept the presented reasoning.)

Of course, different audiences come to the table with different commitments and doubts, and we've just seen that those attitudes can affect what pieces of reasoning one can rationally accept. So how dialectically effective an argument is will depend on *who* its audience is. (I imagine this will affect how pragmatically appropriate it is *to give* the argument in different conversational contexts.)

Moore's argument is directed at a skeptic. We've focused on the skeptic who doubts whether our perceptual experiences give us any justification at all for our perceptual beliefs. Clearly Moore's argument is not very dialectically effective against that skeptic. But it should be clear by now why that's so. The skeptic *has doubts* that prevent Moore's argument from rationally persuading him. There's nothing wrong with the justificatory structure the argument articulates, or with Moore's own reasoning. What's wrong is that the skeptic has doubts he ought not to have.

We just discussed *how dialectically effective* Moore's argument will be against a skeptic. A separate question, and a fifth target of evaluation, is how good of a **philosophical response to skepticism** Moore's argument constitutes. Moore had a variety of anti-skeptical ideas, including his claim that we're reasonably more confident that we have hands than we are of the premises in the skeptic's arguments.⁴⁶ But let's just consider Moore's argument from (1) and (2) to (3). Our question is: how satisfying a philosophical response to skepticism does *that argument* constitute?

Nowadays, it's commonly agreed that an adequate *philosophical response* to the skeptic need not be capable of rationally *persuading* the skeptic that the external world exists, or that we have justification to believe it exists.⁴⁷ Nor need it be capable of persuading someone who's seized by skeptical doubts. What it *does* have to do is diagnose and explain the flaws in the skeptic's reasoning. It has to explain away the intuitions that the skeptic draws support from. These are *not* responsibilities that one has as an ordinary believer. The ordinary believer who's never heard the skeptic's arguments—or who's heard them but rationally believes they've got to be flawed somehow—doesn't need to do anything more, before he can believe with justification that the world is the way it looks to him. But they are responsibilities we have when we're doing philosophy. That's *the business* of philosophy: to diagnose and criticize arguments like the skeptic's.

Clearly Moore's argument, by itself, does little to discharge those responsibilities. I think it *does* offer us a piece of reasoning by which we can acquire justification to believe the external world exists. But it takes a lot of supporting argument—only some of which I've given here—to establish that. If we're to have a satisfying philosophical response to skepticism, it will consist in that supporting argument, not in the reasoning that Moore's argument articulates.

Notes

* Versions of this paper have been circulating since 2001. Many thanks to Robert Audi, Jake Beck, Helen Beebe, Karen Bennett, Alex Byrne, Stew Cohen, Annalisa Coliva, Juan Comesaña, Martin Davies, Greg Epstein, Miguel Fernández, Mark Greenberg, Patrick Hawley, Richard Heck, Tom Kelly, Adam Leite, Christian Piller, Mark Schroeder, Nico Silins, Jonathan Vogel, Ralph Wedgwood, Tim Williamson, Crispin Wright, my spring 2004 graduate seminar, and audiences in Arizona, Mexico City, Paris, Seattle, and Wake Forest, for all their useful comments and feedback.

1. In Moore 1939.
2. Contextualism has come under heavy fire recently: see Sosa 2000; Kornblith 2000; Feldman 1999; Williamson forthcoming; Stanley 2004; Hawthorne 2004, Ch. 2; and Richard 2004.
3. For instance, see Hawthorne 2004, Ch. 1 and forthcoming.
4. See Wright 1985, 2000a, 2002, 2003, 2004; and Davies 1998, 2000, 2003, and 2004.
5. Consider e.g., Plautinga 1993, Ch. 1; and Goldman 1988's notion of "weak justification."
6. For more on these issues, see my 2001, §4.2.
7. Matters would be more complicated if we were talking of justification to *fully believe* P. You can have a degree of justification to believe P that's not yet enough to justify an attitude of full belief. However, I propose to work always with a notion of justification to *partially believe* P: justification to believe P to a

- given degree. When I omit mention of degrees (as I often do), it's just for expository convenience.
8. As I said in the previous footnote, this all needs to be understood in terms of degrees. Additionally, I'm assuming you don't *also* have prima facie justification to believe P from further sources that *aren't* undermined by U.
 9. For more on this relation of antecedent justification, see my 2000. I argue there that it plays an essential role in the most powerful skeptical arguments. See also Klein 1981, §§2.13–15; Klein 1995, n. 16; and Wright's papers cited in note 4, above.
 10. A conservative about H may also count evidence *against* H—that is, evidence that *opposes* the antecedent justification you have to believe H—as undermining the prima facie justification you have (and retain) to believe P. When your antecedent justification to believe H is *undermined*, though, the conservative will claim you cease to meet the conditions necessary *to have* prima facie justification to believe P.
 11. Such a theory assigns H's truth an essential role in *giving* you prima facie justification to believe P; but says that evidence against H has no *intrinsic* undermining potential. It can undermine your prima facie justification for P only indirectly, by raising the likelihood of other hypotheses that *do* have undermining potential. The clearest way for this to obtain would be that *you don't recognize* the role H plays in justifying P. If *you do* recognize that role, then acquiring evidence against H should *at least* justify you in believing *you don't have justification* to believe P. I think that negative higher-order justification does undermine the first-order justification you get from H's really obtaining; but this is controversial. We'll discuss the question in section 6, below.
 12. Compare BonJour's discussion of "background conditions" for a priori justification in BonJour 1998, pp. 126ff. and 137.
 13. Suppose you *do* understand the proof in our example, but you acquire misleading evidence that you don't. That is, you get evidence against H1. I'm inclined to think this evidence would undermine the justification that your (actual) understanding of the proof gives you. If so, then we should treat H1 in the liberal (and truth-requiring) way. If on the other hand we think the hypothesis that not-H1 has no intrinsic undermining potential, then we should treat H1 in the merely truth-requiring way.
 14. Hypotheses can undermine without being non-perceiving hypotheses. For example, consider the hypothesis that Moore *tends to* hallucinate hands, and the hypothesis that his visual experiences *aren't generally reliable*. It's compatible with these hypotheses that, on some particular occasion, Moore is perceiving his hands. So they're not non-perceiving hypotheses. But plausibly, if Moore were to acquire evidence for them, it would undermine his perceptual justification to believe he has hands. Perhaps that's because evidence for them *makes it more likely* that Moore is not perceiving right now. In any event, I will for simplicity confine my discussion to just those undermining hypotheses that are also non-perceiving hypotheses.
 15. BonJour 1985 and Cohen 2002 are good examples. Note that one can be conservative about perception without being conservative across the board.
 16. Burge, Peacocke, Davies, and Dretske also prefer the terms 'entitlement' and 'warrant' to 'justification.' I insist on 'justification.' It ought not to have the

undesired associations they hear it to have. Plus, ‘entitlement’ has bad associations of its own, and ‘warrant’ is already used in too many different ways.

17. The higher-order questions are most prominent in Wright’s 2004.
18. See, e.g., Wright 1985, pp. 435–8 and Wright 2002, pp. 336–8.
19. See Wright’s 1985, pp. 449ff.; 2000a, pp. 152–3, 156–7; 2000b, pp. 212–13; 2003, pp. 66, 68; and 2004. Compare Cohen’s 1999 and 2000, which claim that certain skeptical hypotheses are a priori irrational, so we’re entitled to reject them without evidence.
20. I understand Martin Davies to be a lapsed conservative. He used to express sympathy for a view like Wright’s, which said that Moore’s perceptual justification to believe he has hands rests against “background assumptions” that Moore needs antecedent “entitlement” for (see Davies 2000, p. 401). Davies stresses that normally one won’t *believe* those background assumptions; it’s enough merely to *not doubt* them. But it does seem that you need *justification* to believe a proposition, for it to play the role of a background assumption. Davies characterizes background assumptions as propositions B where the epistemic probability of H for you, given some evidence E, is $p(H|E\&B)$ (p. 396). B can only play that role when it itself has some epistemic probability for you.

In his most recent paper (Davies 2004), Davies drops the requirement that Moore have “positive entitlement” to any background assumptions. His new view is near, or the same as, the view I favor.

21. See, e.g., Hume’s *Treatise* I.IV.2; Reid’s *Inquiry* V.7, VI.20; Strawson 1985; and Dretske 2000.
22. This includes Pollock 1986, pp. 142–8; Brewer 1999; and Peacocke 2004. The details of their accounts are quite different.
23. In Pryor 2000, n. 37, I called this the “phenomenal force” of perceptual experience, thinking of it on analogy with the assertoric force of a public utterance. See also Heck 2000, pp. 508–9; Thau 2002, §2.3; and Burge 2003, p. 543.
24. I imagine there *could* be subjects who represent that the police have arrived in the phenomenologically distinctive way I described. They would seem to ascertain that the police have arrived. But in my example, you do not.
25. That’s somewhat over-stated. I think your experiences also justify you in believing that *you are ascertaining* those propositions, and that *you’re having* those experiences. See note 33, below.
26. Compare McLaughlin’s criterion for question-beggingness in McLaughlin 2000, pp. 104–5. That criterion should probably be amended to require that the subject *recognize* she has justification for the premise only if the conclusion is true. Even so amended, though, McLaughlin’s criterion would still wrongly count the argument (4)–(7), and the argument in the next note, as question-begging.
27. Davies 1998, p. 352 also makes this point.
Here’s another example. I attend to my occurrent mental life, and observe:

(i) I am conscious.

from which I conclude:

(ii) So somebody is conscious.

- Here too, the truth of the conclusion seems a necessary condition for having the justification I do to believe the premise. Yet once again, the argument sounds perfectly respectable.
28. The reference to *what kind* of justification you purport to have for the premises is essential. Suppose your conclusion is “I ingested no memory-altering drugs recently” and your premises concern your activities over the past few days. If your grounds for believing those premises consist of other people’s present testimony, then the case is unremarkable. Evidence that you *did* ingest memory-altering drugs may oppose, but it wouldn’t undermine your grounds for believing the premises. However, if your grounds consist of *your memories* of the past few days, then your grounds would be undermined by evidence that you ingested memory-altering drugs recently. In this second case, we have an example of the type of epistemic dependence we’re considering.
 29. As with Type 4 arguments, the reference to a *kind* of justification for the premise is essential. Arguments may transmit some kinds of justification for their premises but fail to transmit others. See Wright 2000a, p. 141 and Wright 2003, p. 58.
 30. Some will claim that you genuinely do feel cold, but only for a second. I think that’s incorrect. Of course there’s *some* qualitative difference between your entire phenomenology when you think you’re having a sensation of cold, and when you *realize* you’re having one of heat. But this difference is a difference in how you *recognize* or classify your sensation. When you decide that your initial classification was wrong, it doesn’t feel like you’re trying to keep up with a sensation whose intrinsic quality is changing. You feel like you’re recognizing the true character of a sensation that has remained the same. See Feldman 2003, pp. 55–6.
 31. Thanks to Jonathan Vogel for help constructing this example.
 32. It’s unclear to me what Wright’s view is. Clearly he thinks Type 5 dependence suffices for transmission-failure. On the simplest reading, he thinks transmission-failure *coincides with* Type 5 dependence. Matters are complicated, though, in several ways.

(1) Wright sometimes seems to be saying that Type 4 dependence also suffices for transmission-failure. (For example, in some passages on pp. 342–3 of his 2002. These passages also occur in his 2000a and 2003.) However, when he says that he *seems* also to be thinking that Type 4 dependence entails Type 5 dependence. I hesitate because in the crucial passages he shifts to talking about higher-order justification. He *may* think that what transmission-failure coincides with is your needing antecedent justification for the conclusion in order to have justification to believe *that you have justification* for the premises.

(2) Wright talks about “information-dependent” and “inferential” warrants, and his characterization of these makes them sound like justification that relies on antecedent justification to believe something else. Wright *definitely* thinks arguments can exhibit transmission-failure even when your justification for their premises isn’t “information-dependent” (see his 2002, pp. 344ff. and 2003, pp. 60–63). This suggests that we can have transmission-failure even when your justification to believe the premise *doesn’t* require antecedent justification to believe the conclusion (or anything else). But there’s also evidence that Wright understands “information-dependence” more narrowly than my relation of needing

antecedent justification to believe something else, and that he thinks all cases of transmission-failure *do* involve what I'd call Type 5 dependence. (See e.g., Wright 2002, p. 346 and 2000b, pp. 212–13.)

Davies is also difficult to interpret. He first says we have transmission-failure when your justification for an argument's premises "presupposes" or requires a "prior commitment" to the argument's conclusion, or when the conclusion's truth is a "precondition" of your justification for the premises (see his 1998, pp. 350–1, 354). He doesn't specify what those relations amount to. On one interpretation, he too should be read as saying that transmission-failure coincides with Type 5 dependence. But other interpretations are also possible.

Later Davies says we have transmission-failure when your justification for the premises has "background assumptions" that you could not rationally accept while doubting the conclusion (Davies 2000, pp. 402ff.). We'll discuss questions about what you can rationally accept while doubting below.

Davies' most recent account (Davies 2004) comes close to my own. He now grants that Moore's argument *can* transmit justification to its conclusion. He just thinks it cannot "settle the question" whether that conclusion is true. "Settling the question" is supposed to be analogous to convincing a doubter (see Davies 2003, pp. 41–3). So Moore's failure to settle the question is only a dialectical or persuasive failing.

Alston 1986 and Bergmann forthcoming also think that Type 4 arguments can give us justification to believe their conclusion. They do regard them as "epistemically circular" arguments, but they say in this case the circularity is benign. A Type 4 argument's weakness is just that it cannot persuade doubters.

33. As I mentioned in note 25, above, I think your experiences will *also* justify you in believing that *you're seeing* the wall to be red. You're visually ascertaining that *it is* red. From that premise you can infer the simpler conclusion that the wall is not lit by tricky red lights. (If it were so lit, then even if it were red, you wouldn't be ascertaining its redness.) I think this inference is just as good as the inference I cite in the text.
34. Cohen 2002, p. 314. See also Wright 2002, p. 342. The next few sections give the start of a response to Cohen's argument. My full response requires the non-classical system of epistemic probability I set out in Pryor forthcoming. Cohen also gives a second argument, concerning "bootstrapping." I'll have to respond to that argument elsewhere.
35. Compare Jackson 1987, Ch. 6 on the "teasing-out" role of argument.
36. When I talk of a subject's "trying" to believe what he has justification to believe, or "recognizing" that some premises support his conclusion, I'm not imagining him to have explicit attitudes. In normal cases, these will just be tacit commitments of his reasoning, commitments that are manifested in how ready he is to revise or retract that reasoning in response to different sorts of defeating evidence. There are difficult and important issues here, but we can't pursue them now.
37. Sometimes it can be *practically rational* to deliberate with, and act on, attitudes from which you haven't managed to purge all such flaws. (See Harman 1986, Ch. 2.) Nonetheless, believing in this way still constitutes a rational failing.

Relations like the one I'm calling "rational commitment" have been much discussed in moral philosophy. They play a central role in recent work by Broome and Dancy. (See Broome 1999; Broome 2003; and Dancy 2000, Ch. 2–3.) Broome and Dancy explicate the idea in terms of "wide-scope oughts," whose normative force attach to a conditional or disjunctive attitude, rather than to a particular attitude. As Broome would put it, you ought to: *believe Johnny can fly only if you believe someone can fly*. From that, and the fact that you actually do believe Johnny can fly (perhaps without reason), it does not follow that you ought to believe someone can fly: $O(Bp \supset Bq)$ and Bp don't entail OBq . This explication is formally neat, and it makes the non-detachability of the ought vivid. However, Mark Schroeder has persuaded me that it shoulders controversial commitments not required by the core idea of a non-detachable relation between states (some of his argument is set out in Schroeder). These include a questionable semantics for the English auxiliary 'ought,' and predictions of more symmetries than our intuitions validate. The explication I give in the text avoids some of those further commitments.

38. Perhaps *pathological* doubts wouldn't have that power. Suppose you have a nagging belief that you're a brain in a vat, which you recognize to be unjustified, but which you just can't get rid of. But you go ahead and form perceptual beliefs on the basis of your experiences, just like everyone else. Then you'd be exhibiting a kind of irrationality. But in this case it's not clear we should attribute any irrationality to your perceptual beliefs. They might arguably count as well-founded.
39. Like me, Bergmann forthcoming also thinks that having doubts about Moore's conclusion will take away your doxastic justification for believing Moore's premise. Bergmann doesn't discuss what effects doubts have on propositional justification.
40. My account of agnosticism requires the non-classical system of epistemic probability I mentioned in note 34.
41. Wright 1985, p. 347, my italics. See also Davies 1998, p. 351; and Davies 2000, p. 404.
42. Compare Harman's Principle of Positive Undermining in Harman 1986. See also Goldman's account of undermining in Goldman 1986, Ch. 4–5.
43. I hesitate to say that *the argument* justifies you in believing its conclusion—after all, it's flawed. Perhaps we should say your present *understanding* of the argument justifies you. It *does* at least seem to justify you in the higher-order belief that the argument gives you justification for its conclusion. In the same way that I allowed negative higher-order justification the power *to undermine* first-order justification, I'm tempted to also allow *positive* higher-order justification the power *to supply* first-order justification. (Broome defends a practical analogue of this; see his 1999 and 2001, §4.) Notice that if we say this, then we need to posit some opacity about where one's justification is coming from. The subject will think it comes from the argument itself. In the case we're considering, it doesn't.
44. In Pryor 2000, I thought that we should not count a priori skeptical arguments as introducing "defeating evidence" (p. 354). On the current proposal, though, they do. (This *isn't* how the skeptic conceives what he's doing; see below.)

45. Wright's definition of a "coherent argument," in Wright 2002, p. 331 equates:

- (i) being an argument that could move someone *to rationally overcome doubt* about its conclusion

with:

- (ii) being an argument that could move someone *to rational conviction* in its conclusion.

See also Davies 2000, pp. 388, 397ff.. Our discussion in sections 5–6 shows that an argument can have virtue (ii) without having virtue (i). Davies now agrees: see his 2004. Beebe 2001 similarly distinguishes between arguments that "fail to convince" in the sense of not persuading someone who doubts their conclusion, and arguments that "fail to convince" in the sense of not giving one justification for their conclusion.

46. See Moore 1909 and Moore 1918/19.

47. See, e.g., Cohen 1988; DeRose 1995, §17; Pryor 2000; and Byrne 2004.

References

- Alston, William. 1986. Epistemic circularity. *Philosophy and Phenomenological Research*, **47**.
- Beebe, Helen. 2001. Transfer of warrant, begging the question, and semantic externalism. *Philosophical Quarterly*, **51**, 356–74.
- Bergmann, Michael. forthcoming. Epistemic circularity: malignant and benign. *Philosophy and Phenomenological Research*. Available online: <http://philosophy.rutgers.edu/EVENTS/EPIS2003/Bergmann.pdf>.
- BonJour, Laurence. 1985. *The Structure of Empirical Knowledge*. Cambridge, Mass.: Harvard University Press.
- BonJour, Laurence. 1998. *In Defense of Pure Reason*. Cambridge: Cambridge University Press.
- Brewer, Bill. 1999. *Perception and Reason*. Oxford: Oxford University Press.
- Broome, John. 1999. Normative requirements. *Ratio*, **12**, 398–419.
- Broome, John. 2001. Normative practical reasoning. *Aristotelian Society Supplement*, **75**, 175–93.
- Broome, John. 2003. Practical reasoning. In: Bermúdez, José Luis, & Millar, Alan (eds.), *Reason and Nature*. Oxford: Oxford University Press.
- Burge, Tyler. 2003. Perceptual entitlement. *Philosophy and Phenomenological Research*, **67**, 503–48.
- Byrne, Alex. 2004. How hard are the skeptical paradoxes? *Nous*, **38**, 299–325.
- Cohen, Stewart. 1988. How to be a fallibilist. *Philosophical Perspectives*, **2**, 91–123.
- Cohen, Stewart. 1999. Contextualism, skepticism, and the structure of reasons. *Philosophical Perspectives*, **13**, 57–90.
- Cohen, Stewart. 2000. Contextualism and skepticism. *Philosophical Issues*, **10**, 94–107.
- Cohen, Stewart. 2002. Basic knowledge and the problem of easy knowledge. *Philosophy and Phenomenological Research*, **65**, 309–29.
- Dancy, Jonathan. 2000. *Practical Reality*. New York: Oxford University Press.
- Davies, Martin. 1998. Externalism, architecturalism, and epistemic warrant. Pages 321–61 of Wright, Crispin, Smith, Michael, & Macdonald, Cynthia (eds.), *Knowing Our Own Minds: Essays in Self-Knowledge*. Oxford: Oxford University Press.
- Davies, Martin. 2000. Externalism and armchair knowledge. Pages 384–414 of: Boghossian, Paul, & Peacocke, Christopher (eds.), *New Essays on the A Priori*. Oxford: Oxford University Press.

- Davies, Martin. 2003. The problem of armchair knowledge. Pages 23–55 of: Nuccetelli, Susana (ed.), *New Essays on Semantic Externalism and Self-Knowledge*. Cambridge, Mass.: MIT Press.
- Davies, Martin. 2004. Epistemic entitlement, warrant transmission, and easy knowledge. *Aristotelian Society Supplement*, **78**, 213–45.
- DeRose, Keith. 1995. Solving the skeptical problem. *Philosophical Review*, **104**, 1–52.
- Dretske, Fred. 2000. Entitlement: epistemic rights without epistemic duties? *Philosophy and Phenomenological Research*, **60**, 591–606.
- Feldman, Richard. 1999. Contextualism and skepticism. *Philosophical Perspectives*, **13**, 91–114.
- Feldman, Richard. 2003. *Epistemology*. Upper Saddle River: Prentice-Hall.
- Goldman, Alvin. 1986. *Epistemology and Cognition*. Cambridge, Mass.: Harvard University Press.
- Goldman, Alvin. 1988. Strong and weak justification. *Philosophical Perspectives*, **2**, 51–69.
- Harman, Gilbert. 1986. *Change in View*. Cambridge, Mass.: MIT Press.
- Hawthorne, John. 2004. *Knowledge and Lotteries*. Oxford: Oxford University Press.
- Hawthorne, John. forthcoming. The case for closure. In: Steup, Matthias, & Sosa, Ernest (eds.), *Contemporary Debates in Epistemology*. Oxford: Blackwell.
- Heck, Richard G., Jr. 2000. Non-conceptual content and the 'space of reasons'. *Philosophical Review*, **109**, 483–523.
- Jackson, Frank. 1987. *Conditionals*. Oxford: Blackwell.
- Klein, Peter. 1981. *Certainty: a refutation of skepticism*. Minneapolis: University of Minnesota Press.
- Klein, Peter. 1995. Skepticism and closure: why the evil genius argument fails. *Philosophical Topics*, **23**, 213–36.
- Kornblith, Hilary. 2000. The contextualist evasion of epistemology. *Philosophical Issues*, **10**, 24–32.
- McLaughlin, Brian. 2000. Self-knowledge, externalism, and skepticism. *Aristotelian Society Supplement*, **74**, 93–117.
- Moore, G. E. 1909. Hume's philosophy. *The New Quarterly*.
- Moore, G. E. 1918/19. Some judgments of perception. *Proceedings of the Aristotelian Society*, **19**.
- Moore, G. E. 1939. Proof of an external world. *Proceedings of the British Academy*, **25**, 273–300.
- Peacocke, Christopher. 2004. *The Realm of Reason*. New York: Oxford University Press.
- Plantinga, Alvin. 1993. *Warrant: The Current Debate*. New York: Oxford University Press.
- Pollock, John. 1986. *Contemporary Theories of Knowledge*. Totowa, NJ: Rowman & Littlefield.
- Pryor, James. 2000. The skeptic and the dogmatist. *Noûs*, **34**, 517–49.
- Pryor, James. 2001. Highlights of recent epistemology. *British Journal for the Philosophy of Science*, **52**, 1–30.
- Pryor, James. forthcoming. Uncertainty and undermining.
- Richard, Mark. 2004. Contextualism and relativism. *Philosophical Studies*, 215–42.
- Schroeder, Mark. forthcoming. The scope of instrumental reason. Available online: <http://www.princeton.edu/~mschroed/papers.html>.
- Sosa, Ernest. 2000. Skepticism and contextualism. *Philosophical Issues*, **10**, 1–18.
- Stanley, Jason. 2004. On the linguistic basis for contextualism. *Philosophical Studies*, 119–46.
- Strawson, P. F. 1985. *Skepticism and Naturalism: Some Varieties*. New York: Columbia University Press.
- Thau, Michael. 2002. *Consciousness and Cognition*. Oxford: Oxford University Press.
- Williamson, Timothy. forthcoming. Knowledge and skepticism. In: Jackson, Frank, & Smith, Michael (eds.), *The Oxford Handbook of Analytic Philosophy*. Oxford: Oxford University Press.
- Wright, Crispin. 1985. Facts and certainty. *Proceedings of the British Academy*, **71**, 429–72.
- Wright, Crispin. 2000a. Cogency and question-begging: some reflections on McKinsey's paradox and Putnam's proof. *Philosophical Issues*, **10**, 140–63.

- Wright, Crispin. 2000b. Replies. *Philosophical Issues*, **10**, 201–19.
- Wright, Crispin. 2002. (Anti-)sceptics simple and subtle: Moore and McDowell. *Philosophy and Phenomenological Research*, **65**, 330–48.
- Wright, Crispin. 2003. Some reflections on the acquisition of warrant by inference. Pages 57–77 of: Nuccetelli, Susana (ed.), *New Essays on Semantic Externalism and Self-Knowledge*. Cambridge, Mass.: MIT Press.
- Wright, Crispin. 2004. Warrant for nothing (and foundations for free)? *Aristotelian Society Supplement*.