Is Moore's Argument an Example of Transmission-Failure?

James Pryor
Harvard University
<jpryor@fas.harvard.edu>

Draft 2 — 8/12/01

I

Consider the following well-worn example, first put forward by Fred Dretske. You're at the zoo, and in the pen in front of you is a striped horse-like animal. The sign on the pen says "Zebra." Assuming that animal really *is* a zebra, it would seem that your evidence is perfectly adequate to enable you to *know* that it's a zebra. So you know:

(ZEBRA-1) That animal is a zebra.

Now what about the claim that this animal is just a mule painted to look like a zebra? You know that if the animal is a zebra, it isn't a mule, and *a fortiori* it isn't a cleverly-disguised mule. Hence, you know:

(ZEBRA-2) If that animal is a zebra, it isn't a cleverly-disguised mule.

But are you really in a position to know:

(ZEBRA-3) That animal isn't a cleverly-disguised mule?

You may have *some reason to believe* ZEBRA-3. Zoos don't typically try to fool people like that; they have security systems to keep out pranksters; and so on. But your evidence doesn't seem to be good enough to *know* that the animal in the pen is not a cleverly-disguised mule. You haven't made any special tests, or anything like that. So Dretske thinks you don't know it. But he still wants to say that, as long as the possibility that the animal is a cleverly-disguised mule is not a *relevant* epistemic possibility, you *can* know ZEBRA-1.[1]

---

[1] See Dretske 1970.

It seems like we have a failure of Closure here. You know ZEBRA-1, and you know that it entails ZEBRA-3, but your evidence is not good enough to enable you to know ZEBRA-3. And indeed, Dretske presented the case to help show that there could be failures of Closure. As Dretske construes the example, you're in a position to rule out all the epistemic possibilities that are relevant alternatives to ZEBRA-1, but when we're considering ZEBRA-3, *more* epistemic possibilities are relevant, and you're not in a position to rule out all those additional possibilities. So although your evidence is good enough for you to count as knowing ZEBRA-1, it's not good enough for you to count as knowing ZEBRA-3.

Dretske's treatment of these examples still has some defenders.[2] But many Relevant Alternatives Theorists these days would rather *keep* Closure. "Look," they say, "in any *one* context, the set of which epistemic possibilities are relevant is fixed. Either that set includes the possibility that the animal in the pen is a cleverly-disguised mule, or it doesn't. When we consider ZEBRA-1, the mule-possibility is not likely to jump out at us, so we won't regard it as relevant. But when we consider ZEBRA-3, then the mule-possibility *does* strike us as relevant. But what has happened here is that the context *has changed*. The context has changed because what epistemic possibilities we take seriously has changed. In the old context, the mule-possibility is not a relevant alternative, and so doesn't need to be ruled out; in the new context, it is relevant, and so does need to be ruled out. In no single context do we find any violation of Closure. Either your evidence is good enough to rule out all the alternatives to ZEBRA-1 and ZEBRA-3 that are relevant in that context, or it is not. So either you know *both* ZEBRA-1 and ZEBRA-3, or you know *neither*. When you're trying to see whether Closure holds, you should pick a context and stay with it. Don't allow the context to change mid-argument. That would be akin to equivocating."

This is nowadays the most common line for Relevant Alternatives Theorists to take on Closure.[3]

---

[2] See, e.g., Heller 1999.

[3] Stine 1976, Cohen 1988, DeRose 1995, and Lewis 1996 all argue for views of this sort. In Pryor 2001 §1.1, I distinguish *Relevant Alternatives Theorists* from *Contextualists*. It's really

Crispin Wright and Martin Davies have formulated an interesting new complaint about the argument from ZEBRA-1 to ZEBRA-3. They do not want to raise any doubts about Closure. They allow that you do know both ZEBRA-1 and ZEBRA-3 to be true. What they want to know, though, is whether this argument from ZEBRA-1 is capable of *giving you* justification for believing ZEBRA-3? Or would you *already need justification for ZEBRA-3 to be in place*, in order to be justified in believing the argument's premise ZEBRA-1? Wright and Davies think the latter. Hence, as they put it, the ZEBRA-argument is not capable of "transmitting" the justification you have for believing its premise to its conclusion. It illustrates a "failure of transmission," even if it doesn't illustrate a failure of Closure.[4]

We have a case of "transmission-failure," then, when you have justification for believing the premises of some argument, and those premises entail some conclusion, but the argument is not capable of *giving you* justification for believing that conclusion—at least, not when your justification for the premises is the sort it is.[5] This notion of transmission-failure is basically a new piece of terminology for talking about an old phenomenon: the phenomenon of *begging the question.* The reason why the ZEBRA-

---

only the Contextualists who are in a position to give this argument for Closure. Not all Relevant Alternatives Theorists are Contextualists (and as I argue in that paper, some Contextualists like Cohen ought not to be counted as Relevant Alternatives Theorists). For the purposes of this paper, we can overlook those niceties.

[4] This complaint is developed in Wright 1985, Davies 1998, Wright 2000a, Davies 2000, and Wright forthcoming.

[5] An argument might be capable of transmitting certain kinds of justification and incapable of transmitting others. So far, we've been supposing that your justification for believing ZEBRA-1 comes from your visual experiences as of a striped horse-like animal standing idly in its pen. Wright and Davies say that the ZEBRA-argument will not transmit *this* justification to its conclusion ZEBRA-3. But if you had *other* sorts of justification for believing ZEBRA-1, then the argument might very well transmit *that* justification. For example, suppose your justification for believing ZEBRA-1 is that the animal just brayed in a distinctive way that you know only zebras can. Given *this* kind of justification for believing ZEBRA-1, the ZEBRA-argument would appear to be perfectly in order.

argument doesn't give you any justification for believing its conclusion is that the argument is question-begging. The kind of justification you have for believing its premise requires you to *already* be justified in believing ZEBRA-3, before you're entitled to employ this premise in any arguments. So, naturally, arguments starting with this premise will not do anything to enhance ZEBRA-3's epistemic credentials. If your grounds for believing ZEBRA-1 are the sort we've described, you couldn't use the ZEBRA-argument to *establish* ZEBRA-3.

I agree that transmission-failure is a genuine phenomenon, and worth close study. And I agree that the ZEBRA-argument is a good example of it. However, Wright and Davies think this kind of transmission-failure is also exemplified by Moore's famous "proof":

(MOORE-1) Here is one hand, and here is another.

(MOORE-2) If I have hands, then the external world exists.

(MOORE-3) So, the external world exists.

and by the argument:

(BIV-1) Here is one hand, and here is another.

(BIV-2) If I have hands, then I am not a handless brain in a vat.

(BIV-3) So, I am not a handless brain in a vat.

Here I disagree. I will argue that the charge of transmission-failure is appropriate only in cases with a certain epistemological structure, which I think is absent when we're dealing with basic perceptual judgments like *Here is a hand*. In the arguments MOORE and BIV, I will argue, there is no transmission-failure. The kind of justification our experiences give us for believing *Here is a hand* *does* help make it more reasonable to believe MOORE-3 and BIV-3.

II

Let's begin with some general remarks about justification.

You can have justification for believing a proposition P that you don't in fact believe—either because you haven't considered P, or because you haven't noticed that

you have justification for believing it, or because you're overly cautious, or stubborn, or irrational, or whatever.

When you have justification for believing P, that can justify you in believing further things. When it does, we'll say that your justification for believing those further propositions *rests* in part on your justification for believing P. You can have justification for believing Q that rests in part on justification you have for believing P, even in cases where you don't actually believe P or Q. Hence, you need not have *inferred* Q from P, or *based* a belief that Q upon a belief that P, for these epistemic relations to be in place.

Here's an example. Suppose you look at the gas gauge of your car and form the belief that you're out of gas. One thing that might have happened is that you formed a belief about the gauge, and based your belief about the car on that belief about the gauge. In such a case, it is clear that your belief that the car is out of gas rests on your belief about the gauge. But things needn't have gone that way. You might not have given the gauge any thought. You might have formed your belief about the car directly, without inferring it from premises about the gauge. In such a case your belief about the car will have been *formed* without any inference. But it will still rest, epistemically, upon your justification for believing that the gas gauge says "E." After all, your visual experiences do justify you in believing that the gas gauge says "E," and if you were to lose that justification, you would no longer be justified in believing that the car is out of gas. Or perhaps you didn't form *any* beliefs about your car. Perhaps you're unjustifiably paranoid, and you refuse to believe any of the things your eyes tell you. Your experiences would still give you justification for believing that the gas gauge says "E," and that in turn would justify you in believing that the car is out of gas, even if you didn't form either of these beliefs. Given your evidence, you *ought to* believe that you're out of gas—regardless of whether you *do* believe it.[6]

Let's take some proposition P that you're justified in believing, and construct a graph of all the propositions that its justification rests upon, and all the propositions that *their* justification rests upon, and so on. As follows:

---

[6] In section V, we will look at some complications raised by cases like this one.

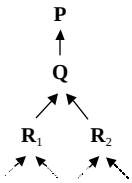

The graph should include not only propositions that you considered and explicitly based your belief in P upon, but *all* supporting propositions where your justification for believing P rests upon your justification for believing those supporting propositions. We say that a proposition like R_1 appears "below" P in this graph when there is a chain of arrows that starts at R_1 and eventually leads to P. Now, if some proposition X never appears below P in this graph, then we say that you have justification for believing P which is *independent of*, or *antecedent to*, any justification you have for believing X.

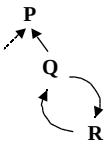

In a case like this, there is a loop between Q and R, so you're not justified in believing either of those propositions antecedently to the other. However you *are* justified in believing Q antecedently to believing P: P doesn't appear "below" Q in your justificatory graph.

Now, suppose you have some argument for Q that employs the premise P, and suppose that Q appears "below" P in your justificatory graph:

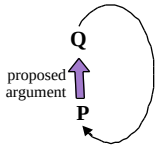

In a case like this, you don't have reasons for believing P that are independent of or antecedent to your reasons for believing Q. So it doesn't seem legitimate here to use P as a premise in an argument for Q. Such an argument would seem to *beg the question* whether Q. If an argument is going to give us justification for believing Q, we ought to be *antecedently justified* in believing that argument's premises.[7]

This is what seems to be going on in the ZEBRA-argument we considered at the beginning. You have visual experiences as of a striped horse-like animal standing in the zoo pen. These seem to justify you in believing that the animal is a zebra. But they're not

---

[7] In Pryor 2000, I distinguished between skeptical principles which say "To know things on the basis of perception, you need to know you're not being deceived by an evil demon," and stronger skeptical principles, which say "To know things on the basis of perception, you need to *antecedently* know you're not being deceived by an evil demon." I argued that only skeptical arguments that employ the stronger principles pose any serious threat. For more on the role that epistemic priority relations play in skeptical arguments, see: Wright 1985, pp. 433, 435-8; Wright 1991, pp. 99-100??; Wright forthcoming, pp. 9-10??; Sosa 1988, pp. 158-9; Klein 1981, §§2.13-15; and Klein 1995, n.16.

really enough, by themselves, to justify you in believing that. They *only* justify you in believing that the animal is a zebra *if* you have some independent or antecedent justification for believing that the animal is not a cleverly-disguised mule. This is why it begs the question to employ the premise that the animal is a zebra in an argument for the claim that it's not a cleverly-disguised mule. You need to be justified in believing it's not a disguised mule, in the first place, before your visual experiences justify you in believing that premise.

I propose that whenever an argument begs the question, or exhibits transmission-failure, it will be because the premises and conclusion stand in this kind of epistemic relation to each other. It will be because the kind of justification you have for believing the argument's premises requires you to have antecedent justification for believing its conclusion.[8] (Later, we will consider whether the class of question-begging arguments is in fact *broader* than this. I don't think it is; I think this proposal does correctly identify the class of question-begging arguments. But we'll get to that in due course.)

If this is right, then our question about whether the arguments MOORE and BIV beg the question will reduce to the question whether, in order for your experiences to justify you in believing *Here is a hand*, you need to be antecedently justified in believing that the external world exists, and that you're not a brain in a vat.

III

There are three ways one might treat general claims like these in one's epistemology:

---

[8] Wright and Davies employ a variety of phrases that seem to express the same idea. See the papers by Wright cited in the previous footnote, and see also: Wright 2000a, pp. 141, 143, 146, 155-6; Wright forthcoming, pp. 2, 3, 5, 7; Davies 1998, pp. 350, 351-2; Davies 2000, pp. 402ff. and 410ff.

In this paper, I am focusing solely on the kind of question-begging relations that arise concerning an argument's *premises*, our entitlement to accept them, their relation to the argument's conclusion, and so on. There is another kind of question-beggingness that can arise, which has more to do with our entitlement to accept *rules of inference*. I will not address that kind of question-beggingness here.

- The external world exists
- I'm not a brain in a vat
- Perceptual conditions are normal (no tricky lighting, etc.)
- My senses are reliable

Let N be some general claim of that sort. The most **conservative** treatment says that, in order to be justified in believing anything about your surroundings on the basis of perception, you need to have independent or antecedent justification for believing N.

A **liberal** treatment, on the other hand, says that for your experiences to justify you in believing things about your surroundings, it only has to be the case that you *lack* evidence for believing that N is false. You don't *also* need to have some positive, antecedent justification for believing that N is true. Nor does N have to actually *be* true. So long as you *lack* reasons for believing that N is false, your experiences are able to give you justification for your perceptual beliefs.

An **intermediate** view *does* require N to be true, but *doesn't* require you to have any antecedent justification for believing that N is true. So long as N *is* true, and you lack reasons for believing that it's false, your experiences are able to give you justification for your perceptual beliefs.

The liberal view and the intermediate view have it in common that:

- you're not required to have antecedent justification for believing that N is true
- but you are required to *lack* evidence that N is false; if you acquire evidence that N is false, that will *defeat* the *prima facie* justification your experiences give you for your perceptual beliefs

I've presented these three views as views about the epistemology of perception; but they generalize. For different beliefs, and different choices of N, different views may be appropriate. Here are two examples.

Suppose you're considering some proof of the Pythagorean Theorem. Let U be the claim that you *understand* and *can follow* the proof. Now, for you to be justified in believing the Theorem, U does have to be true. But you don't need to have *evidence that* U is true. It's the proof itself which justifies you in believing the Pythagorean Theorem. U is just some condition that *enables* that to happen. It is not itself one of the premises

that your justification for believing the Theorem has to rest upon—not even a suppressed, background premise. So the right view about U seems to be the intermediate view.[9]

A second example. You have visual experiences as of a hand. Let R be the claim that your visual experiences are reliable. A *reliabilist* will take R here to have the same status that U had in the previous example: it has to *be true*, for your experiences to justify you in believing you have a hand, but you do not need to have any *evidence or justification for believing* that it's true. *Internalists*, on the other hand, will deny that R needs to be true, for your experiences to justify you in believing that you have a hand. Either they'll take a liberal line, and say that it's enough if you *lack* reason to believe that your experiences are *unreliable*. Or they'll take a conservative line, and say that you do need to have positive, antecedent justification for believing that your experiences are reliable. But in neither case does *the truth* of R make an epistemic difference. It's only your *epistemic situation concerning R* which is important.

So we can see that one might want to handle different cases differently.

Wright and Davies (and many other philosophers) take a conservative line on the epistemology of perception. This underlies their views about transmission-failure. They think that MOORE and BIV exhibit transmission-failure because they think that, just as with ZEBRA, your experiences give you justification for believing the argument's first premise only insofar as you're *antecedently* justified in believing the argument's conclusion. That is why these arguments can't help confirm or make their conclusions any more likely.

Wright and Davies do make allowances that make their views a bit "softer" than other conservative views. Every conservative view makes it a precondition, for a subject's experiences to justify her in believing things like *Here is a hand*, that she be justified in believing certain general background assumptions like *The external world exists* and *I am not a brain in a vat*. But Wright and Davies go on to say:

---

[9] See BonJour's discussion of "background conditions" for *a priori* justification in BonJour 1998, pp. 126ff. and 137.

(Allowance-1) It needn't be the case that the subject is *aware* of this precondition, or that she pays those general background assumptions any special attention.

(Allowance-2) The subject's justification for believing those general background assumptions need not be something she did anything to *acquire* or *earn*. She need not have any *evidence* for them, or anything in the way of a *justifying argument* she could give for them. Instead, she might have some kind of *default entitlement* to believe them. Perhaps there are *a priori* reasons for believing them, which can justify the subject in believing them even if she doesn't know what those reasons are.[10] Or perhaps the subject is justified in believing the background assumptions because they constitute "hinge propositions" or "presuppositions of our epistemic project."[11] The details will not be important to us here.

Though Wright and Davies make these allowances, I still count their views as conservative. Subjects may not often *form* beliefs in the general background assumptions, but on Wright and Davies's views they are *justified* in doing so,[12] antecedently to

---

[10] See Wright 2000a, pp. 152-3 and 156-7; Wright 2000b, pp. 212-13; and Wright forthcoming, p. 17. Similarly, Cohen thinks that certain skeptical hypotheses are *a priori* irrational, and that we're entitled to reject them without evidence: see Cohen 1988 and Cohen 1999. See also BonJour 1985, §§8.3-8.4.

[11] See Wright 1985, pp. 449ff. and Davies 1998, pp. 350ff. and 354.

[12] Wright and Davies make a distinction between "entitlement," on the one hand, and "justification" or "warrant" on the other hand, which I am purposefully glossing over here. As they use these terms, "entitlement" is a more primitive epistemic status, that one does not have to do anything to acquire. (Their notion of "entitlement" is reminiscent of Burge's notion in Burge 1993 and Burge 1996, though it seems to me that there are also important differences. Peacocke?? also distinguishes between "entitlement" and "justification.") I agree that there can be positive epistemic statuses of the sort they have in mind, but I use the term "justification" in such a way that it includes them.

The term Wright contrasts to "entitlement" is "warrant." He articulated a very specific (and idiosyncratic) notion of "warrant" in his Wright 1991; but he nowhere says he intends to be using "warrant"

believing anything on the basis of perception—and they are *required* to be antecedently justified in believing those background assumptions, for their experiences to give them any justification for believing particular truths about their surroundings. That is why Wright and Davies regard arguments like MOORE and BIV as question-begging.

I myself *don't* think those general background assumptions play the epistemological role that Wright and Davies assign them. On my view, your perceptual experiences justify you in believing propositions like *Here is a hand* ***without*** your needing to have antecedent justification for believing general claims like *There is an external world*, or *I am not a brain in a vat*, or anything else of that sort.[13] You don't even need the default, background kind of justification that Wright and Davies describe. I think that, so long as you *lack* any reason to believe that you *are* a brain in a vat, etc. your experiences will justify you in believing *Here is a hand*. Of course, if you acquire some evidence *that you are* a brain in a vat, that will defeat the *prima facie* justification your experiences give you.[14]

---

in the current discussion in the same way he used it there. So in the current discussion, I think we can fairly take his "warrants" to be what most epistemologists would call "justification."

[13] However, I do think that our justification for believing *That animal is a zebra* does require us to have antecedent justification for believing *That animal is not a cleverly disguised mule*; that is why I am willing to count the ZEBRA-argument as an example of transmission-failure. In the terminology of Pryor 2000, the difference is that *Here is a hand* is a *perceptually basic* proposition. *That animal is a zebra*, on the other hand, goes beyond what is really represented by your experiences. (If it turned out that the animal in the pen *is* a cleverly-disguised mule, or a fur-covered robot, we wouldn't say that you've mis-seen it. The error wasn't in what vision reported to you, but in what you went on to believe.) Because of this, when you believe *That animal is a zebra*, you do need some antecedent justification for discounting the possibility that it's a cleverly-disguised mule.

It is quite difficult to tell what propositions are perceptually basic. I believe that *Here is a hand* is perceptually basic, but this choice of example is not crucial. If you don't regard it as perceptually basic, just substitute some other proposition which is.

[14] Perhaps starting to *believe* that you are a brain in a vat, even without evidence, would also defeat your perceptual justification. We will take up that possibility in section V, below.

Let me sketch in a few details.

I think our perceptual experiences give us justification for believing a variety of things. When you have an experience as of P, it gives you introspective justification for believing *that you have that experience*. It also gives you *prima facie* justification for believing that P is true. And it gives you *prima facie* justification for believing things like *I perceive that P*.

What's crucial here is that the justification your experience gives you for believing P does not rest upon any *premises about* what experiences you have, or how reliable your experiences are, or anything like that. Neither does it rest on premises of the form *I perceive that P*. You *may have* justification for believing these other claims, but your perceptual justification for believing P doesn't rest upon it. It's just that often your justification for believing these various claims will come from the same experiential source.[15]

On my view, then, the justification your experience gives you for believing P is in place so long as *you have* that experience, independently of whether *you believe* that you

---

[15] Perhaps nothing *could give us* the kind of justification our experiences give us for believing P, without *also* giving us justification for believing *I have an experience as of P* and *I perceive that P*. If so, that would be an interesting epistemological fact; but it would not show that our justification for believing P rests upon our justification for those other claims.

We habitually say things like: "You based your belief on your experiences," or "You based your belief on the fact that things looked that way." Doesn't this suggest that your perceptual beliefs rest upon premises about what kinds of experiences you're having? It need not. In Pryor 2001 §3.2 I argued that we need a notion of basing even when we're dealing with immediate justification. Suppose you're considering some mathematical claim T that seems obvious to you, and you also have a medium's testimony in support of T. We'll say: if you're rational, you'll believe T because it's obvious, not because of the medium's testimony. We're not suggesting that you base your belief in T upon some psychological premise about how T appears to you. We're urging you to let your belief be controlled and supported by the *immediate* justification you get from considering the proposition, rather than by the testimonial evidence you also possess. I think that talk about "basing your belief on your experiences" works similarly. Here too I think we are talking about letting your belief be controlled and supported by some *immediate* justification your experiences give you. I don't think your perceptual beliefs need to rest on any premises about what kinds of experiences you're having.

do. Because this justification for believing P is only *prima facie*, there are certain kinds of facts that would *defeat* it. However, it's not a condition for *having* the *prima facie* justification that you *first* have justification for believing that those defeating conditions are absent.

I think this liberal view of perceptual justification fits many of our pre-theoretic intuitions about what it's reasonable to believe when. It's only as we become epistemologically more sophisticated that we start to think that the liberal view can't work. However, I think many of the supposed difficulties can in fact be met. As I see it, this liberal view is driven by its naïve appeal. The main work for systematic epistemology is to defend the view against challenges.[16]

IV

Let's survey a number of different ways in which the premises of an argument can epistemically depend upon the argument's conclusion. We want to know what kinds of dependencies would render the argument question-begging and illegitimate.

We'll say that an argument exhibits a **Type-I dependency** just in case your grounds for believing the argument's premises give you sufficient justification for believing those premises *only if* the argument's conclusion is true. Sometimes this kind of dependency has been thought to render an argument question-begging.[17] But I think a little reflection will show that it doesn't. Consider the argument:

(1) Hmmm, I could not have the belief that I exist without that belief's being true.

(2) That means I am *justified* in having that belief.

---

[16] See Pryor 2000 for a contribution to that effort.

[17] See McLaughlin's Principle QB in McLaughlin 2000, pp. 104-5. The principle should probably be amended so that the subject *knows* that her grounds justify her in believing the premises only if the conclusion is true. Even so amended, though, McLaughlin's principle would still be vulnerable to the counter-examples that I give next.

(3) So, someone is justified in believing something.

This argument doesn't seem to be question-begging. It seems to be a perfectly good argument for its conclusion. Yet, (3)'s being true *is* a necessary condition for the reasoning captured in (1) to give me any justification for believing (2). Nothing can justify me in believing (2) unless something justifies *someone* in believing *something*. So this argument exhibits a Type-I dependency. But it doesn't seem to be question-begging.[18]

Here's another example. I attend to my thoughts, and notice:

(4) I think that Maine is north of Massachusetts.

from which I conclude:

(5) So, somebody has thoughts.

Here, too, the argument seems perfectly legitimate. Yet again, the truth of the conclusion does seem to be a necessary condition for me to have the justification I have for believing the premise.[19]

Let's consider a different kind of epistemic dependency. Suppose that an argument's conclusion C is such that its negation would be a *potential defeater* of the justification you have for believing some of the argument's premises. In other words, were you to acquire evidence that not-C, that would *defeat* the kind of justification you have for believing the premises. Of course, whenever you have evidence against a consequence of some set of premises, that tells evidentially against those premises. But I have something more specific in mind here. I'm thinking of cases where evidence that not-C would defeat *the specific kinds* of grounds G you have for believing one (or more)

---

[18] Davies 1998, p. 253, and Wright 2000a, p. 149 also make this point.

[19] Davies 1998, pp. 354-5 uses an example like this, but to make a slightly different point.

Some philosophers think that the argument (4)–(5) *is* question-begging, at least in the context of debates with eliminativists about belief. (See Sainsbury 2000; also Devitt / Boghossian debate??) Personally, I think (4)–(5) is a good argument even in the context of debates with eliminativists. (Is it *my* fault they chose to argue for an obviously false view?)

of the premises. Whether this is so will not follow just from the fact that C is a consequence of a set of premises containing P. It will depend upon the specific nature of C and G, and the interaction between them. For example, suppose C is *I ingested no memory-affecting drugs recently* and your reasons for believing this include premises about your activities over the past few days. If your grounds G for believing those premises consist of other people's testimony, then the case is unremarkable. However, if your grounds consist of *your memories* of the past few days, then those grounds would be defeated, were you to acquire evidence that you *had* ingested memory-affecting drugs recently. When we have an argument whose conclusion stands in this relation to some of its premises, we'll say that the argument exhibits a **Type-II dependency**. The kind of justification you have for believing the argument's premises requires you to *lack* justification for believing that the argument's conclusion is false.[20]

Would this kind of dependency render an argument question-begging? We'll take up that question in just a moment. First, let me introduce a third, and even stronger kind of dependency. This is one we've already encountered. An argument exhibits a **Type-III dependency** when the justification you have for believing the argument's premises requires you to have positive, *antecedent justification* for believing the argument's conclusion. That is, having justification for believing the conclusion is *a precondition* for being justified in believing the argument's premises in the way you do. (Other people might believe the premises on *other* grounds, that don't have the same precondition.)

Type-I dependencies are clearly too weak to render an argument question-begging. Type-III dependencies are clearly strong enough to render an argument question-begging. Earlier in this paper I proposed that an argument is question-begging just in case your justification for believing its premises requires you to have antecedent justification for believing its conclusion. That is, I proposed that question-begging

---

[20] Notice that Type-I dependencies do not entail Type-II dependencies. C might be a necessary condition for your grounds G to justify you in believing P. But you may not know that G is what justifies you in believing P; or you may not recognize that G justifies you in believing P only when C is true. In either case, it's unclear why evidence against C would *have to* render your belief in P on the basis of G less justified.

arguments are just arguments with Type-III dependencies. Is that right? Or ought we to count arguments with Type-II dependencies as also being question-begging?

Conservatives don't acknowledge the existence of any Type-II dependencies which aren't *also* Type-III dependencies. But suppose we don't take a conservative view. Suppose we allow for Type-II dependencies without Type-III dependencies. We allow that a fact D might be a potential defeater of your justification for believing P, without you needing to be antecedently justified in believing not-D, before you're justified in believing P. It's enough, we think, if you merely *lack evidence that D is true*. What should we say about such cases? Would it then be illegitimate to employ P in any argument whose conclusion was not-D?

I do not think these arguments are illegitimate or beg any questions. In the next section, I will try to persuade you of this. But for the moment, let me make a more modest claim. This is that there is an *epistemologically significant difference* between arguments like ZEBRA, on the one hand, and arguments like MOORE and BIV, as non-conservatives understand those arguments, on the other. We should all agree that arguments with Type-III dependencies, like ZEBRA, cannot give anyone justification for believing their conclusions. But on a non-conservative epistemology, arguments like MOORE and BIV do not have the same epistemic structure. Being perceptually justified in believing some animal is a zebra *does* require one to be antecedently justified in believing that it is not a disguised mule; but being perceptually justified in believing *Here is a hand* does *not* require one to be antecedently justified in believing that the external world exists, or that one is not a brain in a vat. So the reasons that render ZEBRA incapable of transmitting justification don't also apply to MOORE and BIV. It may be that the kinds of epistemic dependencies that MOORE and BIV *do* have—Type-II dependencies—are also a kind of epistemic vice. In the next section, I will try to persuade you that that is not so. But even if I fail, and someone comes up with a story to convince us that arguments with Type-II dependencies *are* question-begging, after all, it's important to realize that *some such story is needed*. We can't automatically assume that arguments like MOORE and BIV are question-begging just because arguments like ZEBRA are.[21]

---

[21] Occasionally, it is unclear whether Wright thinks that certain arguments have the kind of epistemic

As I said, I am inclined to accept MOORE and BIV as perfectly legitimate arguments. Or, to take a different example, suppose I look at the wall and it looks red to me. So I believe:

(RED-1) The wall is red.

Now, I know that if the wall is red, it is not white, and *a fortiori* it isn't white but lit by tricky lighting, of a sort that would make it look red:

(RED-2) If the wall is red, it isn't white but lit by red lighting.

Hence I can conclude:

(RED-3) The wall isn't white but lit by red lighting.

I regard this as being a legitimate argument, too. I don't think that my perceptual justification for believing RED-1 requires me to have antecedent justification for believing anything like RED-3.

If I were to acquire evidence that the wall *is* white but lit by red lighting, it would be unreasonable for me to stick to my guns and say, "Well I can see that the wall is red, and from that it follows that the wall isn't white but lit by red lighting." Evidence that the wall is lit by red lighting *would defeat* the justification I have for believing RED-1. But that only shows that the argument exhibits a Type-II dependency. So far, we've heard no story about why that should render the argument incapable of giving us justification for believing its conclusion, when evidence that the wall is lit by red lighting is absent.

---

structure I've called a "Type-III dependency," or whether he's arguing that, *even if they don't*, they're still question-begging. That is, whether he thinks that even arguments with Type-II dependencies are unable to transmit justification to their conclusions. See, e.g., Wright 2000a, pp. 153-5, and Wright forthcoming, pp. 13-14. But primarily, he seems to be employing the first strategy. The dominant picture one gets from his writings is that whenever we have a Type-II dependency, there would also be a Type-III dependency. (E.g., things he says immediately following the passages I just cited indicate that he's thinking of the relevant arguments as having Type-III dependencies.)

V

We have to be careful. There are ambiguities in our casual epistemological talk that could mislead us, or obscure the issues we're considering.

The first of these ambiguities occurs in our talk about what one is entitled to "discount" or "take for granted." Suppose we say you need to be entitled to "discount" the possibility that you're a brain in a vat, or that you need to be entitled to "take it for granted" that you're not a brain in a vat. What do these claims mean?

One thing they could mean is that (i) you need to be justified in believing that you're not a brain in a vat—at least in one of the default, background ways that Wright and Davies discuss. A different thing they could mean is that (ii) you're not required to *bother about* the question whether you're a brain in a vat, until evidence that *you are* a brain in a vat should arise. The difference between these is subtle, but it is important for the issues we're examining. My liberal view can allow that you're entitled to "take it for granted" that the external world exists, and that you're not a brain in a vat, and so on, in sense (ii). But it denies that you need to be able to "take these things for granted" in sense (i). You don't need to have any antecedent justification for believing them. Wright and Davies, on the other hand, think that *you do* need to be able to, and that *you can*, "take those things for granted" in sense (i). If we don't keep these two senses straight, it could cause trouble.

Other ambiguities come into play when we talk about "doubting" an argument's conclusion. This could mean a number of different things.

One kind of "doubt" we can have about a proposition is merely *hypothetical*. This is what goes on when we read Descartes' First Meditation. We don't really abandon or suspend any of our beliefs about our surroundings; we just *entertain the possibility* that those beliefs are false, and think about what follows from the fact that that *is* a possibility.

Is entertaining hypothetical doubts about our perceptual beliefs enough to defeat or annul the justification our experiences give us for those beliefs? Some writers give the impression that it is. For instance, when Wright is discussing the MOORE-argument, he says:

> 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.[22]

Those who hold Relevant Alternatives Theories of the sort we looked at in section I will probably be tempted to agree with Wright here. What possibilities you "seriously entertain" might well affect the range of relevant alternatives that your justification has to answer to.

I on the other hand don't think that *merely entertaining* a possibility by itself has any epistemic force. If you *entertain* the possibility that there is no external world, or that you've ingested hallucinogenic drugs, or anything else of that sort, I think all that does is *raise the prospect* of your perceptual grounds being defeated. It doesn't by itself have any defeating power. No more than raising the prospect of breaking your leg would by itself impair your ability to run.

So I want to set aside merely hypothetical doubts, and just concentrate on *real* doubts. These will give us trouble enough.

What is a "real doubt"? People tend to mean different things when they speak of this. Sometimes "real doubt" is a matter of adopting a certain *psychological attitude*: disbelief, say, or at least the suspension of belief. Sometimes, "real doubt" requires that the doubt be backed up by *evidence*; other times not. I will speak as follows. I will count believing not-P as "doubting that P." (Talk about "doubting that P" suggests that you're *not yet certain* that not-P; but we can ignore that for our purposes.) I will also include as "doubting that P" states of mind where you think P is more likely to be false than true, but you're not yet confident enough that it will be false to believe not-P. And I will include states of mind where, even though you don't (or don't fully) believe not-P, you irrationally withhold from believing P in the face of good evidence. Merely having an open mind as to P, and no (undefeated) evidence either way, will *not* count as having any doubts about P.

---

[22] Wright 1985, p. 437, my italics. See also Davies 1998, p. 351, and Davies 2000, p. 404.

You need not have any evidence against P, to have doubts about it, in the sense I've explained. Not every doubt is a justified doubt. If you *recognize* that your doubt is unjustified, but you can't help doubting anyway, I'll say you have a "pathological doubt." Otherwise your doubt is non-pathological. Notice that your doubt can *be* unjustified without your recognizing that it is. So a doubt can be non-pathological but still unjustified. When you have *justification* for disbelieving or withholding belief in P, I'll say you have "reason to doubt P"—regardless of whether you actually do doubt that P.

Now, I think we should all agree that having *reasons* for doubt can affect what you're justified in believing. Let's go back to the argument RED. If you have evidence that the wall you're looking at is lit by tricky lighting, that will defeat the justification your experiences give you for believing that the wall is red. This is so even if your evidence isn't good enough to justify you in believing the wall *is* lit by tricky lighting. Perhaps it merely makes it somewhat likely that there is tricky lighting. I think that so long as you have *some* positive reasons for doubt about the lighting, that's enough to defeat the justification your experiences give you. (Of course, defeat like justification is a matter of degree; and defeating evidence can in turn itself be defeated. But we will pass over these complications.) In the ordinary case, we don't have any such positive reasons for doubt. Ordinarily, we don't have any evidence whatsoever about tricky lighting.[23]

So reasons for doubt can affect what you're justified in believing. What about *doubts themselves*, independently of whether you have reasons for them? Suppose *you believe* that the lighting is tricky, though you have no evidence for this belief. Would that be enough to defeat your perceptual justification for believing that the wall is red?

To answer this, we need to keep a firm grasp on two epistemological contrasts.

---

[23] You might say, "Well, in that case, then what it's rational for us to do is to suspend judgment about the lighting." Yes, but having an open mind about the lighting is not the same as having doubts about the lighting. There is a difference between not having any evidence whatsoever about the lighting, and having positive evidence that the lighting is tricky, that is not yet enough strong enough to warrant *believing* that the lighting is tricky. I'm only counting the latter as "a reason for doubt." Only it has the power to defeat your perceptual justification for believing that the wall is red. I think that when you *lack any evidence either way* about the lighting, your experiences do justify the belief that the wall is red.

The first contrast is between what you have *justification* for believing, and what you are *rationally committed to believing*, given the beliefs and doubts you in fact already have. For instance, let's suppose you believe that Johnny can fly. This belief *rationally commits you* to having certain further beliefs, and lacking others. If you are not *justified* in believing that Johnny can fly, then you may not be *justified* in those further attitudes. But you would still be rationally committed to them, so long as you maintain your belief that Johnny can fly. Given that belief, for example, it would not be *rational* for you to believe that no one can fly—even if you have plenty of evidence, and are justified in believing, that no one can fly. That's a belief that you can only *rationally* have if you give up your belief about Johnny. So what attitudes it would be rational for you to have, given the beliefs and doubts you already have, need not always be *justified* attitudes. For the beliefs and doubts you already have might themselves be unjustified.

The second contrast is between having *justification for believing* something, and having *a belief that is justified* or well-founded.[24] I do not think that unjustified doubts do have any defeating effect on what propositions you're justified in believing. But for your beliefs to be well-founded, it's not enough that they be beliefs in propositions that you are justified in believing. Your beliefs also have to be *based on* that justification, and they have to be *rational* beliefs. Suppose you believe that P, on the basis of what are in fact good reasons for believing P. But you also have certain doubts that, all things considered, make your belief in P irrational. As we've just seen, those doubts need not be justified, to have this effect. If your belief in P *is* irrational, then it will not be a justified or well-founded belief. So this is a way in which even unjustified doubts *can* have a defeating effect on your beliefs.[25]

---

[24] I take the term "well-founded" from Feldman and Conee 1985. See Pryor 2001, §3.1-3.2 for a bit more on this contrast.

[25] Compare Goldman's account of "undermining" in Goldman 1986, Ch. 4-5.

In light of what I've said about the defeating power of doubts, I should amend something I said in Pryor 2000. I said there that we should not count *a priori* skeptical arguments as introducing "defeating evidence" (p. 354). That's OK, so far as it goes. But we have to appreciate that a skeptical argument can cause a subject to have doubts about whether his experiences justify his perceptual beliefs. As we've just

Perhaps *pathological* doubts would not have this power. Suppose you have a nagging belief that you're a brain in a vat, which you recognize to be unjustified, but you just can't get rid of it. Suppose you also form perceptual beliefs on the basis of your experiences, like everybody else. Then you would be suffering a kind of irrationality. But it doesn't seem right to attribute the irrationality to your perceptual beliefs. It's your belief that you're a brain in a vat that is irrational.

But doubts can be unjustified without being pathological. They can be unjustified without your recognizing that they are. Suppose you have some such unjustified, but non-pathological, doubt about a proposition N such that, were you to have *reason* to doubt N, that would defeat the justification your experiences give you for your perceptual beliefs. For example, suppose you suspect that your color vision might not be working properly. This doubt is in fact unjustified, but you have not realized that. In such a case, I'm inclined to say that your doubt *would* make it irrational for you to form any beliefs about color, on the basis of your visual experiences. Even though your experiences might very well be giving you justification for those beliefs.

Now let's consider how all this bears on the topic of question-beggingness. Suppose you doubt that some argument justifies you in believing its conclusion—or suppose you believe something that rationally commits you to doubting that. Then 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 does not matter whether your doubt is justified. Even an unjustified doubt about whether the argument gives you justification can make it irrational for you to accept the argument's

---

seen, even if those doubts are unjustified, they can affect what it's *rational* for the subject to believe, and as a result they do have one kind of defeating power. Notice, though, that they will have that power only over people who are taken in by the skeptical argument, and start to doubt whether their perceptual beliefs really are justified. (Perhaps also over people who *ought* rationally to be taken in by the skeptical argument, given their other beliefs, but who pay the skeptical argument no heed.) The skeptical argument will have no epistemic effect on those who haven't heard it, or on those who discern its flaws.

I am grateful to Ralph Wedgwood for pressing me on the contrast between justification and rational commitment; and to Patrick Hawley for pressing me on the defeating power of skeptical arguments.

conclusion. Hence, the argument will not be able to *rationally persuade you*—at least, not until you give up the relevant doubt. But this need be no fault of *the argument's*. Suppose that the argument in question is a perfectly respectable proof of the Pythagorean Theorem; and that you have unjustified doubts about whether the premises of this proof are true, or about whether your mathematical reasoning abilities are working properly. If you are *in fact* justified in believing the premises, and you are *in fact* able to understand and follow the proof, then the proof *will* give you justification for believing its conclusion. You won't be in a position to *rationally accept* that conclusion, not until you adjust the other attitudes you hold. But *the proof* is perfectly respectable. The fault lies with *you*, for having doubts you have no good reason to have.

Let's say that an argument is **dialectically effective** for a subject just in case the subject could rationally accept the conclusion of that argument, on the basis of the argument. The example we just considered teaches us that dialectical effectiveness is no sure measure of the quality of an argument. To be sure, some arguments are dialectically ineffective because they're bad arguments. But other arguments are dialectically ineffective, for certain subjects, because of unjustified attitudes that those subjects hold. Reflecting on what those subjects can or cannot rationally believe will not be a good way to measure what an argument provides them with *justification* for believing.

This is an important lesson.

When people are trying to determine whether arguments for some conclusion Q are question-begging, they sometimes ask questions like this:

- If I came to have doubts about Q, would this argument help allay, or *enable me to rationally overcome*, those doubts?[26]
- Would it be possible to *rationally combine* the justification I have for believing these premises with doubt about Q?[27]

If what I've said so far is right, then these tests for question-beggingness are subtly flawed. For suppose we're dealing with an argument like MOORE or BIV or RED, which

---

[26] See Wright forthcoming, p. 2; and Davies 2000, *passim*.

[27] See Davies 2000, pp. 397-9, 401.

exhibits only Type-II dependencies. So, if a subject *had reasons* for doubting the argument's conclusion, that would defeat her justification for believing the premises; but she doesn't need to be antecedently justified in believing the conclusion, to be justified in believing the premises. Now what if the subject just *happens to* doubt the conclusion, without having any reasons for doing so? From what we've seen, this will make it irrational for her to believe the argument's premises, and hence, irrational for her to accept the argument's conclusion (at least, on the basis of that argument). So the argument won't be capable of "rationally overcoming" her doubt. It will be dialectically ineffective for any subject who already has doubts about its conclusion.

But why should that reflect poorly on the argument? We've already seen that unjustified doubts can make it irrational for one to accept the conclusions of *perfectly good* arguments. Here we're just dealing with a special case, where the doubts happen to be about the very proposition which is the argument's conclusion. But I don't see why that should make any fundamental epistemological difference. As before, the argument might very well be *giving* the subject justification for believing its conclusion. In light of her doubts, it is not rational for her to accept that conclusion. But that's not the argument's fault. It's the subject's fault, for having doubts where no doubt is justified.

Hence, I see no reason to regard arguments with Type-II dependencies as being question-begging, or as having any *epistemological* vice. To be sure, they will be *dialectically* ineffective for some subjects. But this is also true of arguments we all acknowledge to be perfectly respectable. When you *don't* have any positive reasons for doubting the conclusions of these arguments, then the Type-II dependency ought not to make any difference. The arguments will transmit whatever justification you have for believing their premises to their conclusions.[28]

---

[28] Beebee 2001 also distinguishes between arguments that "fail to convince" in the sense that they could not persuade someone who doubts their conclusions, and arguments that "fail to convince" in the stronger sense of never transmitting justification to their conclusions. Unlike me, however, she counts arguments that "fail to convince" in the first sense as question-begging. This is because, following Jackson 1987 Ch. 6, she takes begging the question to be essentially a dialectical phenomenon. I think that is a mistake. Begging the question is an informal *fallacy*; it ought therefore to be a *bad thing to do*, no matter *who* one's audience is. The criterion I've proposed for begging the question respects that observation.

VI

At this point, I've disarmed a number of complaints that conservatives might lodge against arguments like MOORE and BIV and RED.

"Look," they'll say, "in order for you to be justified in believing *Here is a hand* on the basis of your experiences, you have to take it for granted that you're not a brain in a vat, and so on—and you have to be *entitled* to take those things for granted."

That is so, I reply, in one sense of "entitled to take it for granted." It has to be the case that you don't need to bother about the possibility that you're a brain in a vat, when there is no evidence that you are one. But that does not show the argument to be question-begging. You don't need to be antecedently entitled or *justified in believing* that you're not a brain in a vat.

"Look," they'll say, "if you were confronted with evidence that you *are* a brain in a vat, you couldn't rely on your perceptual experiences to establish that you're not."

That is so, I reply, but it only shows that the argument in question exhibits a Type-II dependency. Evidence that you're a brain in a vat defeats the justification your experiences give you for believing *Here is a hand*. That does not show that, even when evidence that you're a brain in a vat is *lacking*, the justification your experiences give you for believing you have hands can not contribute to your justification for believing that you're not a brain in a vat.

"Look," they'll say, "arguments like these won't help anyone to *rationally overcome doubt* about their conclusions."

Jackson and Beebee's notion of "begging the question" does not. On their account, an argument that "begs the question" for one audience need have nothing wrong with it when employed in other dialectical settings. And as I've been arguing, arguments that "beg the question" in their sense might be *perfectly good arguments*, despite their dialectical limitations. They might *give justification* for believing their conclusions to many subjects for whom they're dialectically ineffective. For these reasons, it seems unnatural to me to use the term "begging the question" in the way that Jackson and Beebee do.

That may be so, I reply, but it does not show the arguments to be question-begging, either. Arguments like MOORE and BIV and RED might be dialectically ineffective against opponents who are inclined to doubt their conclusions. But we've seen that dialectical ineffectiveness is no sure measure of the quality of an argument. To know whether an argument is *really* question-begging, we have to consider not what it would be *rational* for one to believe, but what the argument *gives one justification* for believing. And I've argued that, when you *lack any evidence against* the conclusions of these arguments, you *do* have perceptual justification for believing their premises, and the arguments *do* transmit that justification to their conclusions. So the arguments *give you justification* for believing their conclusions. If you've got unjustified doubts or beliefs that prevent you from rationally accepting those conclusions, that's your problem.

Of course, in all of this discussion I've been helping myself to the assumption that the right epistemology of perception is a liberal one. And that assumption can certainly be challenged. I'd like to close by looking at one of the challenges.

VII

Stewart Cohen has complained that views like mine would validate a certain kind of "epistemic bootstrapping" that we think is not really legitimate. Suppose I have experiences as of a red wall. According to liberals about perception, these experiences can justify me in believing that the wall is red, without me first having to establish that my color vision is reliable. But I can also introspect and determine that I'm having an experience that represents the wall as being red. Hence, it looks like I can conclude that on this occasion my experience is representing the wall correctly. Call that one "trial" of my color vision. A few more "trials" of the same sort, and it looks like I have the beginnings of an inductive argument that my color vision is reliable, after all. But that's crazy, Cohen says. It is not acceptable to employ color vision in this way to establish its own reliability.[29]

---

[29] See Cohen forthcoming. These "bootstrapping" arguments are also discussed in Vogel 2000, §III,

In response, I pose a dilemma. Either (a) I know ahead of time that each of these "trials" can *only* yield the result "The world and my color vision correspond." Or (b) it's an open possibility, for all I know, that a "trial" might sometime reveal my color vision to be inaccurate.

Now, I think if I'm in a situation of sort (a), *that alone* is enough to explain why the "trials" do nothing to justify me in believing my color vision is reliable. *Whenever* you design an experiment that you know can only yield a single conclusion, that is *enough* to render the experiment incapable of enhancing the epistemic credibility of that conclusion. There's no need to bring up further questions about whether we do or don't have antecedent knowledge about how reliable our instruments are.

Consider an example. Suppose I know *for certain* that a given thermometer is a reliable measure of water temperature. It's reliable enough to yield knowledge. But now I want to know how reliable I am at gauging water temperature when I dip my fingers in the water. So I design an experiment in which I first test the water temperature with the thermometer, then I dip my fingers in the water, then I compare the two results. Now, I'm a little clumsy and so I occasionally drop the thermometer into the water. When I do so, I have to reach deep into the water and fish it out. I'm worried that this will affect the experiment, so I decide in advance that whenever I drop the thermometer into the water, that "trial" in my experiment will be aborted.

Now as I'm running my experiment, suppose we discover that, although I sometimes drop the thermometer just by accident, I also have a neural condition that causes me to drop the thermometer whenever its reading is different than my own estimate of the water's temperature would be. As a result, I will never encounter any "trials" where the thermometer's reading and my estimate diverge. Any "trial" that isn't aborted will be one where the thermometer and I agree. So I know ahead of time that my experiment can only deliver one conclusion. This renders the whole experiment useless.

---

and in Fumerton 1995, pp. 173ff.??; though Vogel and Fumerton only use them to criticize reliabilist views. On the face of it, it seems that a reliabilist could avail himself of the same response that I give below.

(Of course, in practice, I might be able to glean some useful information from how often I drop the thermometer and am forced to abort a "trial." But let's set that aside.)

So my experiment is useless. And notice that this verdict has nothing to do with my failing to have reliable access to the water's temperature. I *do* have reliable access to the water's temperature, via the thermometer, and *I know* that I do. The experiment is useless because I know in advance that the only conclusion it will yield is that my own estimate of the water's temperature is as accurate as the thermometer. Regardless of whether they really are equally accurate.

The same thing might be going on when I'm trying to determine how reliable my color vision is. I might know ahead of time that the only result my "trials" can yield is "The world and my color vision correspond." If so, then I think these experiments are useless, for the same reason my "trials" in the water-temperature experiment were useless. There is a difference between the two examples. In the water-temperature case, I had some way of measuring the water-temperature that I antecedently knew to be reliable; in the color-vision case, that is not so. But I don't think that difference makes any difference here. We already have a good explanation of why the experiment does not give us justification for believing its conclusion. There is no cause to lay blame on the liberal assumption that our experiences give us justification even when we don't antecedently know them to be reliable.

That's what I want to say about case (a), where I know ahead of time that I'm only going to get one sort of result. What about case (b), where it is an open possibility that a "trial" will report that my color vision *fails* to track the real color of the wall? In those cases, I think, a series of positive "trials" *would* give me justification, albeit maybe very slight justification, for believing that my color vision is reliable. After all, disconfirming evidence *could* have turned up, yet it didn't. The wall may have looked to me to have an undulating color, that I independently know no real wall has. The fact that that *didn't* happen counts to some small degree in support of the claim that my color vision is reliable.

Consider a related example. Philosophers like to joke about using a newspaper to establish its own reliability. In one sort of case, you pick up the newspaper and say, "This newspaper says that P, and you know, P is true—see, the newspaper says so. So on this

occasion at least, the newspaper got things right." In practice, I think the fact that the newspaper says what it does *is* some small evidence that the paper is reliable. Think of the many ways of failing to be reliable that have been ruled out: the newspaper isn't full of random ungrammatical strings, it doesn't follow up its claim that P by saying "…and not P," etc. Of course, you may have *already known*, or assumed, that the paper would not be defective in those ways. If so, then reading the paper won't give you any useful information about the paper's reliability. We would then be in a case of sort (a).

Another sort of case is where you buy a second copy of the same paper, and find that it also says that P.[30] Does this improve your justification for believing that the paper is reliable? Well, at least now you know that the paper doesn't fail to be reliable by containing contrary claims in each copy. Of course, you may have *already known*, or assumed, that the paper would not be defective in that way. If so, then buying the second copy won't give you any useful information about the paper's reliability. We would again be in a case of sort (a).

I think things stand the same way with our color-vision experiment. To the extent that we allow for the possibility of negative results, the fact that I don't encounter them counts to some small degree in support of the claim that my color vision is reliable. If we *already know* that all the results will be positive, then we're in a case of sort (a), and the experiment is useless for that reason.

I don't see anything in this diagnosis that impugns the liberal assumption that our experiences give us justification even when we don't antecedently know them to be reliable. On the liberal view, there is *no general bar* to arguing from premises that are justified by vision to conclusions about how reliable vision is. *The specific line of reasoning* that Cohen criticized may very well be objectionable. Arguments that vision is reliable may need to employ a larger evidence-base, compare different experiences against each other, and take into account holistic considerations like systematicity and simplicity. But if this is so—if the kind of bootstrapping argument that Cohen criticized

---

[30] See Wittgenstein, *Philosophical Investigations* §265.

is illegitimate—that is for the reasons I've outline here, and not because of any fault in liberalism.

Bibliography

Helen Beebee, “Transfer of warrant, begging the question, and semantic externalism,” *Philosophical Quarterly* 51 (2001), 356-74

Laurence BonJour, *The Structure of Empirical Knowledge* (Harvard Univ. Press, Cambridge, Mass., 1985)

Laurence BonJour, *In Defense of Pure Reason* (Cambridge Univ. Press, Cambridge, 1998)

Tyler Burge, “Content preservation,” *Philosophical Review* 102 (1993), 457-88

Tyler Burge, “Our entitlement to self-knowledge,” *Proceedings of the Aristotelian Society* 96 (1996), 91-116

Stewart Cohen, “How to be a fallibilist,” in James Tomberlin, ed. *Philosophical Perspectives* 2 (1988), 91-123

Stewart Cohen, “Contextualism, skepticism, and the structure of reasons,” in James Tomberlin, ed. *Philosophical Perspectives* 13 (1999), 57-90

Stewart Cohen, “Basic knowledge and the problem of easy knowledge,” *Philosophy and Phenomenological Research* (forthcoming)

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, 2000), 384-414

Keith DeRose, “Solving the skeptical problem,” *Philosophical Review* 104 (1995), 1-52

Fred Dretske, “Epistemic operators,” *Journal of Philosophy* 67 (1970), 1007-23

Richard Feldman and Earl Conee, “Evidentialism,” *Philosophical Studies* 48 (1985), 15-34

Richard Fumerton, *Metaepistemology and Skepticism* (Rowman & Littlefield, Lanham, Maryland, 1995)

Alvin Goldman, *Epistemology and Cognition* (Harvard Univ. Press, Cambridge, Mass., 1986)

Mark Heller, "Relevant alternatives and closure," *Australasian Journal of Philosophy* 77 (1999), 196-208

Frank Jackson, *Conditionals* (Blackwell, Oxford, 1987)

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

David Lewis, "Elusive knowledge," *Australasian Journal of Philosophy* 74 (1996), 549-67

Brian McLaughlin, "Self-knowledge, externalism, and skepticism," *Aristotelian Society Supplemental Volume* 74 (2000), 93-117

James Pryor, "The Skeptic and the Dogmatist," *Nous* 34 (2000), 517-49

James Pryor, "Highlights of Recent Epistemology," *British Journal for the Philosophy of Science* 52 (2001), 1-30

R.M. Sainsbury, "Warrant-transmission, defeaters, and disquotation," in Enrique Villanueva, ed. *Philosophical Issues* 10 (2000), 191-200

Ernest Sosa, "Beyond skepticism, to the best of our knowledge," *Mind* 97 (1988), 153-88

G.C. Stine, "Skepticism, relevant alternatives, and deductive closure," *Philosophical Studies* 29 (1976), 249-61

Jonathan Vogel, "Reliabilism leveled," *Journal of Philosophy* 97 (2000), 602-23

Crispin Wright, "Facts and certainty," *Proceedings of the British Academy* 71 (1985), 429-72

Crispin Wright, “Skepticism and dreaming: imploding the demon,” *Mind* 100 (1991), 87-115

Crispin Wright, “Cogency and question-begging: some reflections on McKinsey’s paradox and Putnam’s proof,” in Enrique Villanueva, ed. *Philosophical Issues* 10 (2000a), 140-63

Crispin Wright, “Replies,” in Enrique Villanueva, ed. *Philosophical Issues* 10 (2000b), 201-19

Crispin Wright, “(Anti-)sceptics simple and subtle: Moore and McDowell,” *Philosophy and Phenomenological Research* (forthcoming)