

<recto>

<CN>10

<CT>**When Warrant Transmits***

<CA>James Pryor

<H1>*I.*

We can ask about doxastic warrant—which of your beliefs are reasonable, or epistemically appropriate?—and we can ask about prospective or propositional warrant—what things do you have **warrant to believe**?

You have **all things considered** warrant to believe Q when Q is the right, or epistemically appropriate, thing for you to believe overall. There should be a story about *why* Q is the right thing to believe: what factors warrant or count in favor of believing Q. As theorists we want to be able to discuss such factors, while allowing that they may be **opposed** by other factors that warrant believing alternative hypotheses. Or they may be **undermined** by information that makes them count less strongly in favor of believing Q. In the vocabulary I prefer, we say that the factors at least constitute **prima facie** warrant to believe Q—and leave it open to further discussion how complex bodies of prima facie warrant should add up.

Sometimes your warrant to believe Q includes warrant to believe **some premises** that entail or ampliatively support Q. For example, what warrants you in believing *An animal is in*

* Thanks to audiences at UNAM, CUNY, Rochester, Indiana, Reed, Buffalo, Hong Kong, Sydney, Canberra, Seoul, Boulder, St Andrews, and Brown; and especially to Martin Davies, Earl Conee, Sinan Dogramaci, Yuval Avnir, Nico Silins, Matt Kotzen, Patrick Hawley, Paul Boghossian, and Crispin Wright for their comments.

(And one more generous interlocutor whose comments I can no longer attach a name to.)

the pen may be some visual warrant to believe *That's a zebra in the pen*. Let W be your warrant to believe Q, the factors that count in favor of believing Q. I'll represent the "counting in favor of" relation with a hollow arrow, thusly:

<DIS>

W



Q. An animal is in the pen.

In this case, W consists in *visual warrant to believe a premise P* that entails Q. We can spell that out like this:

<DIS>

W. [Visual warrant to believe P] where P is: That's a zebra in the pen



Q. An animal is in the pen.

It will be useful to be able to switch easily between talking about *the premise P*, and talking about *the warrant W* you have for Q by virtue of having certain warrant to believe P. Sometimes this is done by overloading the interpretation of a label: using it sometimes for warrants, other times for premises. For instance, we might write:

<DIS>

Z. That's a zebra in the pen.



Q. An animal is in the pen.

and alternately use “Z” to designate the premise P, or the warrant W. This shortcut can ease discussion and is common in practice, so it will be useful to continue it here. Context should always settle when we’re talking about a warrant, and when about a proposition. Despite this terminological convention, though, we should not forget the substantive difference between warrants and propositions. *The warrant Z* is not itself a proposition or premise. Neither is *the premise Z* ever a warrant to believe Q: not even when it entails Q.¹ Rather, it’s *the visual* (or whatever) *warrant you have for* premise Z (which we may also use “Z” to designate) that constitutes a warrant to believe Q.

On many views, warrants *don’t always* consist in warrants to believe supporting premises. For example, what warrants you in believing *It looks as though there’s a zebra in the pen* may just be *your visual experiences*, or perhaps *the fact that you have* those experiences. Not any *premises about* those experiences, that you need warrant for believing. Let’s label the warrant you have here E, and the hypothesis it supports LOOKS. I’ll represent such warrant with brackets, like this:

<DIS>

E. [Visual experiences as of a zebra]



LOOKS. It looks as though there’s a zebra in the pen.


In such cases we’ll only use “E” to designate the warrant your experiences give you to believe LOOKS, not to designate any proposition.

Sometimes we’ll be considering factors that warrant belief in a conclusion only *against the background of* other warrants you possess. For example, on many views, your visual experiences won’t be enough by themselves to warrant believing *That’s a zebra in the pen*.

¹ There is a view that denies this; see Silins 2005, §4.2 and his references to Klein in fn. 19.

But they will warrant that belief against the background of warrant R to believe your vision is reliable. I'll represent that like this:

<DIS>



Here E on its own *isn't* warrant to believe Z: only R+E together count in favor of Z. But R has more of a background role than E. So long as you *do possess* warrant R, you might reasonably believe things on the basis of your experiences *without ever considering* or having any beliefs about how reliable your senses are. That is, so long as you have warrant *to* believe your senses are reliable, your perceptual beliefs might not need to be *based on* that warrant. E might on its own be an adequate *basis* for believing Z.

As you'll see, this is not *my own view* of perceptual warrant; but we want to have the apparatus on hand to talk about it.

<H1>II.

Dretske says that reasoning like this:

<DIS>

Zebra-1. [Visual experiences as of a zebra]



Zebra-2. That's a zebra in the pen.

↓ *This arrow represents not a "counting in favor of," but a move in*

reasoning.

Zebra-3. That's not a mule cleverly disguised as a zebra.

fails; and he says that its failure illustrates a violation of Closure. That is, it may be that premise Zebra-2 entails Zebra-3, and you have warrant (namely, Zebra-1) to believe Zebra-2, without you having warrant to believe Zebra-3.

Wright is sympathetic to the claim that Zebra reasoning would be defective: you couldn't reasonably base a belief in Zebra-3 on this kind of inference. But he resists the claim that this needs to be a violation of Closure. If you have warrant to believe Zebra-2, then perhaps you must *also have* warrant to believe Zebra-3. But it doesn't follow that your warrant to believe Zebra-2 *constitutes* your warrant to believe Zebra-3. It may rather be that you'd never get to the point of having warrant for Zebra-2 without some other warrant for Zebra-3 antecedently in place. So what we have here isn't a failure of Closure, but rather just a **failure of your warrant to transmit** through the reasoning from Zebra-2 to Zebra-3. Because Zebra-2 isn't itself part of your warrant to believe Zebra-3, you can't reasonably base a belief in Zebra-3 on Zebra-2.²

² I'll be drawing on Wright 2000a, 2002, 2003, 2004, and 2007. Klein 1981 had a preliminary version of this notion ("evidence for" closure, as opposed to "being justified" closure), and Wright began articulating it in his 1985. See also Davies 1998, 2000, 2003a, and McKinsey 2003. These authors contrast failures of Closure to transmission-failure: see, for example, Wright 2000a, pp. 140–1, 143, 157; 2002, pp. 331–2, 335; 2003, pp. 57–8, 61, 67–8. But whether these really do diverge is sensitive to what we mean by "Closure." You do have warrant to believe each of Zebra-2 and Zebra-3; so it's true we don't have any failure of this principle:

Closure for Propositional Warrant. If you have warrant to believe some premises, which you recognize entail Q, then you also have warrant to believe Q.


However, some have argued that Closure is better understood as a principle about doxastic warrant, like:

Closure for Doxastic Warrant. If you have warrant to believe some premises, and correctly

For this to be a useful diagnostic tool, we need a systematic story about when we should expect warrant to transmit, and when we shouldn't. What features of this case are behind our judgment that you can't reasonably base a belief in Zebra-3 on Zebra-2? And if we're to avoid violating Closure, why is it that Zebra-3 so conveniently happens to be independently warranted every time we get to the point of having warrant for Zebra-2?

Here's one natural story. As we said at the end of the previous section, perhaps your visual experiences don't justify belief in Zebra-2 *on their own*, but only against the background of other warrants you possess. So our diagram ought really to look like this:

<DIS>




We might express what bothers us about the inference from Zebra-2 to Zebra-3 by saying “You were already taking Zebra-3 for granted, when you moved to Zebra-2!” So perhaps Zebra-3, or something close to it, is part of the background warrant that's needed to provide your warrant to believe Zebra-2:

deduce Q from them, then the belief in Q you so form will be a warranted belief.

And in the Zebra case, a subject could correctly deduce Zebra-3 from Zebra-2, but *wouldn't* be able to reasonably base a belief in Zebra-3 on that deduction. Arguably, *this* Closure principle will fail whenever warrant-transmission does. (See Silins 2005, §5 for relevant discussion.)

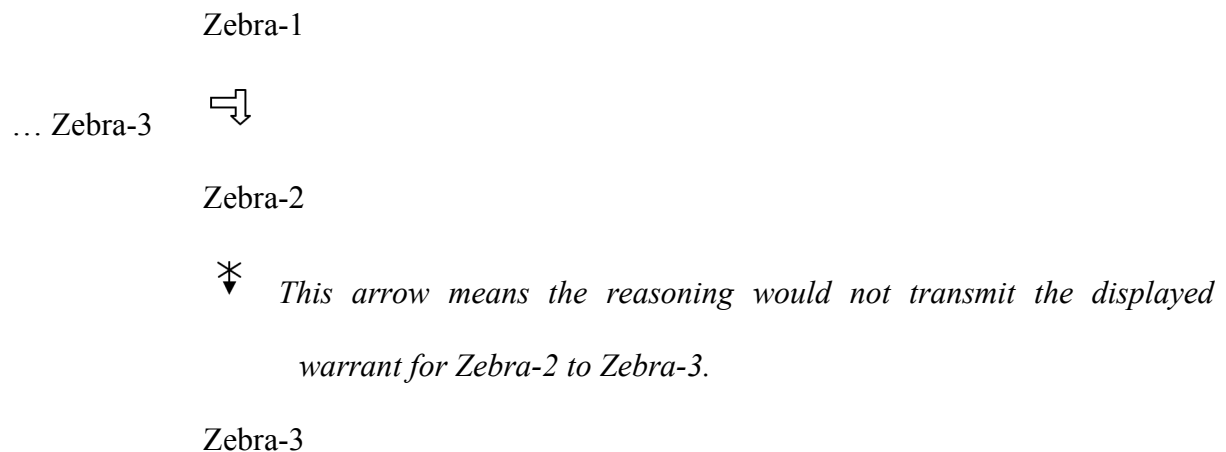
The phenomenon we're discussing is also sensitive to what *kind* of warrant you have for Zebra-2. If your warrant for Zebra-2 wasn't coming from Zebra-1, but rather from the fact that the animal just brayed in the distinctive way that only zebras can, then plausibly you *would* be in a position to eliminate the mule hypothesis, and so *would* be able to infer Zebra-3 from Zebra-2.

<DIS>



Arguably, any case with this structure will bother us in the same way. When Zebra-3 is among the background you need antecedent warrant for, to have the warrant you do to believe Zebra-2, then that warrant for Zebra-2 won't itself further warrant or count in favor of Zebra-3. Consequently, you won't be in a position to reasonably believe Zebra-3 on the basis of Zebra-2. We can call this the **Background Warrant (BW) Model** of transmission-failure:

<DIS>



Some tricky details aside, this is the most plausible and recognizable form of the phenomenon Wright has in mind. Our central two questions in this essay will be: Q1. Does Moorean reasoning, like:

<DIS Proposition>

Moore-1. [Visual experiences as of hands]

Moore-2. I have two hands.

Moore-3. I'm not a disembodied handless spirit, who is merely hallucinating hands.

exemplify the same BW Model? I say it does not. We'll focus even more on: Q2. Can we identify some *other* general model of transmission-failure that Moorean reasoning does exemplify? I will float and criticize some candidates. In the end, I don't think anyone has yet identified a pattern of transmission-failure that we have any consensus Moorean reasoning exemplifies.

We'll begin by sorting out various details.

<H1>III.

Let's say you have an **earned warrant** when the warrant is constituted at least in part by some experience or awareness or insight or understanding that you have. Typical cases of earned warrant will be when you perceive that something is so, or when past observations warrant believing further conclusions. But introspective warrant can be earned, too: the experience of having a headache (or on some views *awareness of* your headache) gives you warrant to believe you have a headache. A priori warrant can be earned, too: your understanding of the relation \neq , the singleton function, and the empty set, together give you warrant to believe $\emptyset \neq \{\emptyset\}$.

We can separate *warranted beliefs* into those which are based on *premises*, or *inference* from other information you possess; and those which are not. We can separate *warrants to believe* into those which are constituted in part by warrant you have to believe other things, and those which are not. These contrasts will not coincide. Recall our earlier example:

<DIS>

E. [Visual experiences as of a zebra]

R ⇨

Z. That's a zebra in the pen.

In this scenario, your belief in *Z* may not be inferred from any premises. Nonetheless, part of the warrant you have for *Z* is the warrant you have to believe *R*, that your senses are reliable. You need warrant *to believe* *R*, though you don't need to *actually believe* or even *consider* *R*, for your belief in *Z* to be warranted.

I'll say that a belief is **inferentially warranted** if it's warranted by virtue of being inferred from other warranted beliefs. Derivatively, we can say that **a warrant is inferential** when it's constituted by *the availability* of some inference to the subject from other, antecedent warrants. On the other hand, a warrant is **mediate** just when it's constituted (at least in part) by warrant you have to believe or doxastically accept other things. What our example illustrates is that a belief can be non-inferentially warranted, and so its warrant be non-inferential, even though that warrant is still mediate. You *might* infer *Z* from *R* and some claim about what experiences you're having; but in the case we're considering, the warrant you have for *Z* isn't constituted by the availability of such an inference. It's constituted by your merely *having* the experiences you do, and also *having* warrant to believe *R*.

A warrant is **immediate** when it's not constituted by warrants to believe or doxastically accept other things. Recall the example:

<DIS>

E. [Visual experiences as of a zebra]

⇩

LOOKS. It looks as though there's a zebra in the pen.

Here, what warrants you in believing LOOKS is just your having certain experiences. Or on some views, it'd be your introspective *awareness of* having them. In neither case is warrant to believe any other claims involved. Warrant of this sort would be immediate.

In addition to earned warrants, some philosophers countenance the possibility of **unearned warrants**. These will be warrants that aren't constituted by any experience or awareness or insight or understanding the subject has. For example, some philosophers talk of us having *default justification* to believe that our senses are reliable, and that we're not brains in vats: these hypotheses, they say, are just intrinsically epistemically probable, without the subject needing to do or experience anything to make them so.³ Some philosophers say that cogito beliefs are warranted just by virtue of being *hyper-reliable*, that is, impossible to form falsely.⁴ And so on.

It's controversial whether there are any unearned warrants.⁵ But many philosophers do countenance some. Wright is among them. He calls unearned warrants "entitlements"; he reserves the term "justification" for earned warrants; and he uses "warrant" for the broad class that includes both earned and unearned varieties. I prefer the term "justification"; I think it fits history and broader usage better to apply that term to *any* kind of warrant to believe or doxastically accept things. If there are entitlements or unearned warrants, I'd call them "justification," too. However, in the present discussion I will stick close to Wright's own usage.⁶

³ See for example White 2004, Davies 2000, Field 2000, Field 2005, Cohen 1988, esp. §V, Cohen 1999, Cohen 2000, Zalabardo 2005, and White 2006.


⁴ See Pryor 2006a and 2006b for discussion.

⁵ Conee and Feldman's "evidentialism" might be understood as the doctrine that all warrant is "earned," in the way we're understanding that here (see their 2004).

⁶ In his 2007, Wright says he uses "warrant" only for all things considered epistemic statuses; but I'm not following him in that regard. I will speak both of all-things-considered and prima facie warrants. Burge 1993

Recall again the example:


<DIS>



I said that your warrant here to believe *Z* is mediate. This will be so *regardless* of whether your warrant to believe *R* is earned or unearned. Even if what warrants you in believing *R* is some default presumption in its favor, rather than any experience or insight you've had, it's still true that part of what warrants you to believe *Z* is warrant to believe something else. So your warrant to believe *Z* is mediate.

On the other hand, some theorists including me think our warrant for some perceptual beliefs is immediate. If *Z* is such a belief, then your warrant would instead be structured like this:

<DIS>



I'm not sure whether claims about zebras are immediately warranted in this way, but I do think that some perceptual beliefs are. At a minimum, our simplest perceptual beliefs, like

and 1996, Dretske 2000, and Peacocke 2004 use the terms "warrant" and/or "entitlement" in ways that roughly parallel Wright's (but there are also differences). Other philosophers use these terms for a variety of different epistemic statuses: for example, Meiland 1980; Pollock 1983; and Plantinga 1993a, 1993b. Wright himself uses the notion of "warrant" differently in his 1991.

There is light ahead, are. I call this view **dogmatism about perception**.

If dogmatism about perception is right, then some instances of Moorean reasoning won't have this structure:

<DIS>

Moore-1. [Visual experiences as of hands]

... ⇨

Moore-2. I have two hands

↓

...

but will instead have this:

<DIS>

Moore-1. [Visual experiences as of hands]

⇩

Moore-2. I have two hands.

↓

...

and so neither Moore-3, nor anything else, will be among the background warrant you need, to have the warrant you do for Moore-2. *There is no* such background warrant. If dogmatism is right, then, these instances of Moorean reasoning won't exhibit the BW Model of transmission-failure. It remains to be seen whether there's some other model of transmission-failure that they do exhibit.

A **Moorean dogmatist** thinks that there isn't—that, whatever its faults, the Moorean

reasoning *is capable* of transmitting warrant to Moore-3. That reasoning *can* be a way to acquire some warrant to believe Moore-3; and it can be reasonable to believe Moore-3 on the basis of it.

I've defended Moorean dogmatism elsewhere.⁷ My preference is for a version of dogmatism where your perceptual warrant comes from the *phenomenology* of your experience, which I think can be shared between veridical perceivers and hallucinators. However, our discussion here will have mostly to do with dogmatism in general: with any view that holds that your perceptual warrant is immediate. This includes reliabilist and disjunctivist versions of the view, too; and views which would explain your perceptual warrant in terms other than what the phenomenology of experience is like.⁸

It's not straightforward *which* perceptual beliefs are good candidates to be immediately warranted. In the Zebra reasoning, I think your warrant *does* plausibly depend on some background assumptions. This is because Zebra-2 goes beyond what's really reported by your experiences. If it *is* a cleverly-disguised mule, or a fur-covered robot, we wouldn't say that you've *misperceived* it. The error wasn't in what you saw, but in what you went on to believe. On the other hand, if you're having illusory experiences as of, say, a light ahead, you *are* victim to a perceptual error. Things aren't the way vision is reporting them. So I'd say that *There is light ahead* is specifically reported by your senses, and is immediately warranted. It's not obvious which example *I have two hands* is more like. But for expository convenience, I'll treat it as a candidate to be immediately warranted, like *There is light ahead*.

Care is needed not to confuse dogmatism with alternative epistemologies.

Dogmatism might naturally be summarized with the slogan "You're allowed to *take it*

⁷ See Pryor 2000 and 2004.

⁸ For a list of dogmatist-like views, see <<http://www.jimpryor.net/research/dogmatist.html>>.

for granted that your senses are reliable, that you're not a brain in a vat, and so on." If you acquire evidence of your *unreliability*, or that you're envatted, that will *undermine* your perceptual warrant; but no antecedent warrant to believe those undermining hypotheses *fail* to obtain is needed. You're "taking it for granted" that those hypotheses are false in the sense of being vulnerable to them as defeaters, without having antecedent warrant to rule them out.⁹

This slogan "You're allowed to take it for granted . . ." can also be read differently. It can be read to mean that you have an *entitlement or warrant to believe* that your senses are reliable, that you're not a brain in a vat, and so on—perhaps a warrant you didn't need to do anything to earn. This view is not dogmatism. It's just one version of the competing view that perceptual warrants are mediate. This is a popular alternative to dogmatism; it's the view Wright espouses, and that Davies held until around 2003.¹⁰ On this view, your perceptual warrants aren't *inferential*. Your experiences as of hands may warrant you in believing you have hands without you needing to *perform* any inferences or even needing to *think about* whether you're a brain in a vat. Neither do you need to have any special experience or insight. The antecedent warrants in question are unearned "entitlements." But to be warranted in your belief that you have hands, on the basis of your hand-like experiences, *you do need to have* antecedent warrant to believe you're not a brain in a vat. It can't come from your perception-based beliefs in the external world; it has to be in place independently.

On this sort of view, the Moorean reasoning exhibits the same BW Model of transmission-failure that we're attributing to Dretske's Zebra reasoning.

I mentioned earlier that there are tricky details with applying the BW Model of transmission-failure to Dretske's example. Let's confront some of them. Consider this alleged part of the Zebra reasoning:

⁹ See here Davies 2004, and Davies 2000, p. 404; see also Burge 2003.

¹⁰ See fn. 3, above, for some other proponents.

<DIS>

Zebra-1

... Zebra-3 ⇩

Zebra-2. That's a zebra in the pen.

Do you *really* need antecedent warrant to believe you're not looking at a disguised mule, to be perceptually warranted in believing Zebra-2? What if that alternative is beyond your ken—e.g. because you've never heard of mules, or because you never would have guessed that animals can be disguised to look like other animals?

One response here would be to say that you can have warrant *to* believe things that you're not yet *able* to believe or even entertain, and to maintain that *you do after all* need warrant *to* believe you're not looking at a disguised mule, even when that alternative is not one you're yet in a position to consider. A different response would be to say that the mule hypothesis is one you'd need antecedent warrant to rule out *when it arises for you* as a question. Prior to that, perhaps you don't. In cases where the subject is proposing to reason from Zebra-2 to Zebra-3, however, the question whether it's a disguised mule has arisen for the subject, so in those cases her warrants will need to have the structure diagrammed above.

Perhaps these responses can be satisfyingly developed. But there is some real doubt whether it's really Zebra-3 that our warrant to believe Zebra-2 depends on. Might it not instead be something more general?¹¹ Perhaps it's rather this:


<DIS Proposition>

Zebra-4. Zoo pens don't often contain animals disguised to look other than they are.

¹¹ See here also Silins 2005, pp. 78–9.

If the structure of the Zebra reasoning is really this:


<DIS>



then the BW Model we've articulated no longer applies.

Perhaps Zebra-4 and Zebra-3 are closely enough related that some minor variant of the BW Model could apply. The idea may be: Zebra-3 is *really* supported by Zebra-4, or by the interaction of Zebra-4 with Zebra-1, and would already need to be so supported for you to have the warrant you do for Zebra-2. So reasoning from Zebra-2 to Zebra-3 wouldn't track the real structure of your warrants.¹² They'd really look like this:

<DIS>



I think it will be extremely difficult to spell out in detail why such structures should obstruct warrant-transmission. In explicit reasoning we very often “overshoot” the real warrants that we're relying on. For instance, suppose I visit a remote plain of Antarctica, and

¹² Compare Wright 2002, pp. 334–5.

see what looks to be a woman ahead. I reason:

<DIS Proposition>

Plain-1. [Visual experiences as of a woman.]

Plain-2. There's a woman walking up ahead.

Plain-3. So I'm not the first person to walk on this plain.

My reasoning here doesn't accurately track the real structure of my warrants. We can undermine my warrant to believe Plain-2—give me evidence that I'm incapable of distinguishing men from woman at this distance—while intuitively leaving in place the same warrant I have for Plain-3. So Plain-2 is really *stronger* than anything my warrant for Plain-3 depends on. I might more accurately have reasoned:

<DIS Proposition>

Plain-1. [Visual experiences as of a woman.]


Plain-4. There's *a person* walking up ahead.

Plain-3. So I'm not the first person to walk on this plain.

Yet I'd be very reluctant to say my actual reasoning, from Plain-2 to Plain-3, is *unreasonable*. I think I *can* acquire warrant to believe Plain-3 by reasoning in that way, and can reasonably base my belief in Plain-3 on such reasoning.

What is the difference between the bad Zebra reasoning and the permissible Plain reasoning? Arguably, both have the same structure:

<DIS>





3

and we want to say your warrant for 3 strictly comes from the interaction of 1 with some of the stuff in the background, like 4, rather than from 2. Yet in the one case reasoning from 2 to 3 can still seem reasonable; in the other case not. What makes the difference? If the BW Model of transmission-failure is to be extended beyond cases where *3 itself* is in the background, these details need to be ironed out.

I don't think it will be easy to do that; I'm not sure it will even be possible. But for our purposes we can set these difficulties aside. If dogmatism is true, then *no* model of transmission-failure which involves background warrants for the 2-premise will apply to Moorean reasoning. For dogmatism says our perceptual warrant for Moore-2 doesn't depend on *any* background warrant.

<H1>IV.

When people articulate their qualms about dogmatism, they often seem to be applying a general principle like this:

<EXT>

No Fallible Immediacy (NFI) If some warrant *W* could be possessed by you in situations other than *A*—for example, if it could be possessed in each of situation *A* and *B*—then *W* can't immediately warrant you in believing that *A* obtains, rather than the alternatives.¹³

¹³ Prima facie, this principle is what seems to be behind arguments like: "Dogmatism has no resources to prevent the conclusion that I can acquire elementary knowledge by vision plus elementary inference of a

But when I articulate this principle, I find that few are willing to embrace it. They're right to hesitate. Many epistemologies on offer nowadays countenance some form of immediate warrant, and often that warrant is thought to be fallible. Can you have immediate but fallible warrant to believe you have a *headache*? If so, then let A be the situation in which you genuinely do have a headache, and B the situation where you don't, but still are warranted in thinking you do, and we have a counterexample to NFI. Can you have immediate but fallible warrant to believe the answer to *a simple logic puzzle* is A? Perhaps you're wrong, the answer is really B. The warrant you have is possessible in situations (such as the real situation) where A is false; yet it's an immediate warrant for A. So we have another counterexample to NFI.

Indeed, Wright himself seems committed to denying NFI. For he thinks we have unearned warrants to believe we're not brains in vats, that we instead have genuine perceptual contact with our surroundings. This warrant is fallible: we might be wrong. But neither does Wright think it's constituted by our warrant to believe other premises. So these unearned warrants can be fallible yet immediate.

If there's no *general* difficulty with fallible yet immediate warrants, we need to know why there should be some *special* difficulty in the case of perception.

Perhaps the problem is that perceptual warrants are earned. Perhaps what's true is not NFI but a more restricted principle:

<EXT>

No Earned Fallible Immediacy (NEFI) If some *earned* warrant. W could be possessed by you in situations other than A, then W can't immediately warrant you in believing that A obtains, rather than the alternatives.

proposition whose being false would predict my enjoyment of exactly the same visual experience. That is absurd." (Wright 2007, §IV.4)

Wright can countenance fallible yet immediate *unearned* warrants, while refusing to allow fallible yet immediate perceptual warrants, because they would be earned.¹⁴

What could motivate this restriction? What keeps it from being ad hoc?

I'm going to suggest a line of reasoning which may be tempting people here. This will be speculative. But I think it's worthwhile articulating the reasoning, and saying what's wrong with it.

Consider a situation in which you already *know* ahead of time what warrant you're going to acquire for A. For example, suppose you already know that when you open your eyes you'll have visual experiences as of an amber wall. In such cases, it's very plausible that *actually having* the experiences shouldn't make any epistemic difference as to whether the wall really is amber. It shouldn't make you any more confident that the wall is amber, rather than beige but lit by tricky orange lighting.

But now suppose you *don't* yet know what warrant you're going to acquire when you open your eyes. *Actually having* experiences as of an amber wall would then tell you you're among those situations where *you will* have these experiences. It's tempting to think *that's all* it can tell you. For, after you factor in the knowledge that you're among the situations where you will have the experiences, what epistemic role is left for the experiences to play? Didn't we just decide, a paragraph ago, that *actually having* the experiences shouldn't make any further difference, beyond what you've now factored in?

And now, haven't we arrived at NEFI? All that your experiences can by themselves warrant you in believing is that you have them. They can't give you an immediate warrant to believe the wall is amber rather than beige but lit by tricky lighting.

There's a subtle mistake in this reasoning. The mistake is to equate "learning that you

¹⁴ Thanks to Carrie Jenkins for identifying this strategy.

will have the experiences” with “learning *only* that you will have the experiences.” On a view where your experiences immediately warrant you in believing the wall is amber, the information that you will have the experiences should support *not just* the claim that you will have them, but also, prospectively, whatever claims those experiences will immediately support. That’s why actually having the experiences makes no difference: if you knew in advance that you’d have them, then you should already have incorporated the warrant they give you to believe the wall is amber. It doesn’t follow that your experiences *only* warrant believing you have those experiences; and hence, only that you’re *either* seeing an amber wall *or* seeing a beige wall lit by tricky lighting. If it’s true that your experiences only warrant that, it needs to be shown another way.

A different source of resistance, to *Moorean* dogmatism in particular, is just that it licenses Moorean reasoning, and that reasoning *sounds bad* to many philosophers. We need some account of why it *should* sound bad. The simplest hypothesis would be that it *really is* bad.

Some philosophers are tempted to diagnose the flaw in Moorean reasoning in probabilistic terms, and argue that this conflicts with what the Moorean dogmatist says. Moore-3 can be formulated in such a way that its negation *entails* you’d have the experiences you do. Then even if your experiences did confirm Moore-2, which entails Moore-3, by orthodox Bayesian thinking those experiences should *disconfirm* Moore-3 rather than supply warrant to believe it. Whenever not-M entails E but M doesn’t, the probability of M given E should be *lower* than the prior probability of M.¹⁵

Some philosophers want to explicate transmission-failure in probabilistic terms like

¹⁵ For a list of philosophers pressing or responding to this argument, see <<http://www.jimpryor.net/research/bayesian.html>>.

these.¹⁶ But I propose to set the probabilistic considerations aside for this discussion. There are many complications at play: What's the relation between probabilistically confirming and supplying warrant? Should the epistemic effect of *having* experiences be the same as conditionalizing on *the premise* that you have them? There's so much to discuss that we need to take it up in other forums.¹⁷ Here I'll suppose that transmission-failure is meant to be a novel diagnostic tool, not just another vocabulary for expressing familiar probabilistic theorems.

The complications aside, I doubt that probabilistic considerations really *could be* the source of our deepest-rooted intuitive resistance to Moorean reasoning. Intuitively, the following reasoning seems just as problematic:

<DIS Proposition>

Moore-1. [Visual experiences as of hands]

Moore-2. I have two hands.

Moore-3*. *I have two hands and I'm not a disembodied handless spirit, who is merely hallucinating hands.*

Yet, whereas Moore-2 merely *entailed* the earlier Moore-3, it's *logically equivalent to* Moore-3*. So if your experiences supply *any* probabilistic confirmation to Moore-2, they have to confirm Moore-3* to the same degree. There's no basis for any *probabilistic* critique of the step from Moore-2 to Moore-3*. Yet our intuitive resistance to this reasoning and to the original Moorean reasoning seem to come from a common source.

Where now are we? The options on the table are: Either (i) we accept Moorean dogmatism, so Moorean reasoning can be reasonable. In that case, we owe an account of why

¹⁶ For a list, see <<http://www.jimpryor.net/research/transmission.html>>.

¹⁷ See Pryor (ms-a).

that reasoning *seems* intuitively to be problematic. Or (ii) we reject dogmatism altogether, as a false epistemology of perception. Instead, we go for something like Wright's epistemology, and say that Moorean reasoning fails because it exemplifies the BW Model of transmission-failure. We might combine with that, or pursue separately, the idea (iii) that there's something probabilistically objectionable about Moorean reasoning. This needs much discussion, and we're going to set it aside here. Or, (iv) we say that even if dogmatism is true, and so Moorean reasoning doesn't exemplify the BW Model of transmission-failure, nonetheless there are *other* (non-probabilistic) models of transmission-failure which it does exemplify. So even if dogmatism is true, Moorean reasoning should still fail.

I'm trying to present dogmatism in a favorable light. But my goal here is not so ambitious as to establish it is true. Rather, my goal is to eliminate option (iv). I think future debate is best focused on the other options. (A fifth, higher-order option will be discussed in §VII, below.)

I'm not sure whether Wright disagrees about this. Often he seems to be backing option (ii); and in 2007 it's hard to read him as entertaining anything else. But as we'll see, there's enough interpretive unclarity that he may sometimes be read as backing option (iv). My goal is to persuade you that that's a bad strategy. If you want to charge Moorean reasoning with transmission-failure, you're best off arguing that dogmatism is false and that Moorean reasoning exhibits the same BW Model of transmission-failure that we're attributing to the Zebra reasoning. (Or getting into the details of what probability dynamics are at work.)

I've proposed a dogmatist defense of Moorean reasoning in my 2004. It will be useful to quickly review what I said there; that will be Section V. In Section VI, we'll take up the interpretive task of sorting out exactly how Wright is conceiving of transmission failure. In Section VII, we'll consider and criticize possible models for transmission failure other than the BW Model.

<H1>V.

If the Moorean dogmatist is to be believed, Moorean reasoning can be reasonable. Why then does it strike us as so intuitively problematic?

To answer this, we need to discuss three or four key questions. The first is: What is the epistemic effect of false higher-order warrant? For example, consider a mathematician who reasons correctly to P, but then acquires warrant to believe, falsely, that the proof strategy she used is flawed. She now has higher-order warrant to believe, falsely, that her original reasoning fails to warrant belief in P. How confident should she be that P? Plausible things can be said on either side of this question; but in the end, I and many others think the mathematician's higher-order warrant should render P somewhat less warranted for her. Warranted beliefs *about* your warrants, even when those beliefs are false, can undermine the prima facie support the lower-order warrants provide.

The second question is: What is the epistemic effect of reasoning that's *in fact* flawed but which *seems* compelling? Suppose our mathematician carefully reasons through another proof, of Q, which *does* have some subtle flaw that she's unaware of. Will she have warrant to believe Q? Perhaps she has inductive warrant: she knows that proofs she finds compelling, and proofs in this journal, have good track-records. But set that aside. Does she have any non-inductive warrant to believe Q? Does the apparent compellingness of the proof, or her experience of seeming to deduce Q in that way, give any warrant for Q?

I don't know what's the right thing to say about this. I can be persuaded either way. What we say about Moore's argument needs to track what we say.

Suppose we say the mathematician *does* have some warrant for Q, despite the flaw in her proof. Then consider a philosopher who's reading skeptical arguments to the effect that

we have no perceptual warrant; or anti-dogmatist arguments to the effect that NEFI is true, and so experiences don't by themselves, immediately, warrant any perceptual beliefs. If dogmatism is true, such arguments are flawed. But the philosopher may not yet see the flaws, or may not see them as flaws. The arguments may seem compelling to her. Then she'll be in a situation like the mathematician's: she may be warranted in believing the skeptic's conclusions, despite the flaws in the arguments. And then, drawing on our answer to the first question, the philosopher's warrant for these false claims *about* her perceptual warrant may well undermine that perceptual warrant. Even though her experiences really do provide it.

That's not to say that the skeptic is *correct*. The skeptic says that *none of us* are perceptually warranted. We're only here describing a mechanism by which *people who read skeptical arguments and find them compelling* may end up in the kind of situation the skeptic falsely claims we're all in. Doing epistemology may be harmful to your epistemic health. Other people, who are ignorant of the skeptic's arguments, or are (permissibly) unmoved by them, can retain their immediate perceptual warrant to believe Moore-2, and can reasonably go on to infer Moore-3 from it.

On the other hand, though, perhaps flawed arguments *can't* warrant belief in their conclusion. In that case, the dogmatist is committed to saying that philosophers *lack* warrant to believe the skeptic's and the anti-dogmatist's conclusions, even when the arguments for those conclusions seem compelling. To sort out what's going on in that case, we need to attend to a third question.

Suppose you merely believe P, without warrant for doing so. You also recognize that P entails Q, and nothing about that entailment serves to undermine your confidence in P. What then is your epistemic relation to Q?

It doesn't seem right to say that you have *warrant* for Q. All that Q has going for it is that it follows from P, an *unwarranted* belief of yours. But neither does it seem enough to say

that Q is just unwarranted, in the way that any other arbitrary unsupported proposition is. Your unwarranted belief in P seems to put some kind of pressure on you to believe Q too. The overall most reasonable course may be for you to believe neither. But if, failing that, you believe P yet resist believing Q, you're manifesting some further epistemic flaw. If you're going to be unreasonable and believe P, then you "ought" in some sense to believe Q too.

This phenomenon is familiar in the ethics and practical reasoning literature. In recent years Broome has discussed how it operates in epistemology too.¹⁸ I've yet to find a fully satisfying vocabulary for describing the phenomenon. Broome began by talking about "normative requirements," as opposed to "reasons": your unwarranted belief in P normatively requires you to believe Q, but since it's unwarranted, it isn't a *reason* for Q. In my 2004, I talked about "rational commitment." Both expressions are non-optimal. For one thing, the general phenomenon here can also be present in cases where the relation between P and Q is merely ampliative; so talk of "requirement" and "commitment" is too strong. For another thing, talk of "commitment" may mislead. We *don't* want to say that the subject is conclusively under any rational obligation to believe Q. Rather, she'd be *best* off believing neither P nor Q.¹⁹

Perhaps it'd be helpful to borrow an idiom from ethics, and call the relation between your belief in P and a prospective belief in Q one of **conditional** or **hypothetical support**. If you really had warrant for P, on the other hand, that warrant would **categorically support** believing Q. Hypothetical support is partly a matter of what such categorical support you *would* have, if you had that warrant for P. But it's also more than that: when your belief in P

¹⁸ See the series of papers beginning with Broome 1999.

¹⁹ Also, one referee thought it more plausible to say it's not your belief in P itself, but rather your doxastic *commitment* to P, which is constituted or manifested by your believing, that "commits" you to any further beliefs.

hypothetically supports believing Q, and you nonetheless resist believing Q, you will thereby be exhibiting some kind of (categorical) epistemic defect.

Some philosophers describe these phenomena in terms of relevance relations that hold between P (or sometimes your belief in P) and *the proposition Q*, rather than a prospective belief in Q. That may be adequate for some cases, but it won't generalize properly. Intuitively, the same phenomenon is at work when a subject believes, without warrant, that on balance her evidence tells equally in favor of Q and not-Q. In such a case, I'd like to say that the subject's unwarranted belief about her evidence hypothetically supports *suspending belief* whether Q. If we were restricted to relations whose right-hand arguments were propositions, we wouldn't be able to say that. (In such cases, the subject's unwarranted belief about her evidence may well hypothetically support *believing* propositions *about her relation to Q*; but it should *also* hypothetically support the prospective attitude of *suspending belief* whether *Q itself* is true.) Similarly, on the left-hand side, if a subject is unwarrantedly agnostic about P, that attitude may hypothetically support being agnostic about further matters. So the phenomenon we're talking about is fundamentally a relation between actual and prospective *attitudes*, not just a relation between propositions.

Here's another generalization of the phenomenon, which will be crucial to our discussion. Suppose you have some warrant W that (categorically but *prima facie*) supports believing Q. This warrant is vulnerable to being undermined by warrant to believe U. Now suppose you *merely do* believe U, without warrant. Earlier, we said that mere belief in P can hypothetically support believing Q. Similarly, I think mere belief in U can **hypothetically undermine** the warrant W gives you to believe Q. In my 2004, I described this by saying your belief in U "rationally obstructs you" from believing Q on the basis of W. You're already manifesting a rational flaw by believing U without warrant; but you'd be manifesting a further flaw if you went ahead and believed Q on the basis of W, while maintaining your

belief in U.

It's clear to me that *this is* a flaw, but unclear how it relates to other, more familiar epistemic statuses. Perhaps having doxastic warrant—*warranted beliefs*, not just warrant *to* believe—requires that your beliefs not be hypothetically undermined in this way. I'm not sure; but this is a proposal I do tentatively defend.²⁰

In any event, let's return to our philosopher who finds flawed skeptical or anti-dogmatist arguments compelling. This philosopher unwarrantedly believes that she's in no position to have perceptual beliefs. She believes that her experiences don't by themselves, immediately, warrant any perceptual beliefs. Or she unwarrantedly believes things that hypothetically support those conclusions. In response to our first question, we said that *warranted* belief in such conclusions should (categorically) undermine the warrant her experiences really do give her. Given what we're saying now, *unwarranted* belief in those conclusions (or in premises that hypothetically support them) should *hypothetically* undermine in a parallel way. For a philosopher with such beliefs, it'd be epistemically defective to believe things just on the basis of her experiences—even if those experiences are in fact giving her categorical warrant *to* so believe.

I've described mechanisms by which philosophers who have certain epistemic views—or beliefs that hypothetically support those views—would have their perceptual warrant for Moore-2 categorically or hypothetically undermined.²¹ Those mechanisms are genuine; but they can foster the illusion that your warrant for Moore-2 is structured like we imagined your warrant for Zebra-2:

²⁰ See Pryor 2004 and Pryor (ms-b). Bergmann 2004 and 2005 shares many important elements of my view; though there are also some differences.

²¹ Note I have not proposed that the *mere possession of* critical faculties, or more conceptual sophistication, makes you worse off. Those would be different sorts of claims. It's only when your faculties *are deployed in the ways I described* that you may make yourself epistemically worse off than necessary.

<DIS>

Moore-1. [Visual experiences as of hands]

... ⇨

Moore-2

That is, they foster the illusion that your warrant for Moore-2 *isn't* immediate, but rather depends on antecedent background warrant to believe things aren't as the skeptic proposes. I want to emphasize we've now seen how these mechanisms may be operating even if that's false, and your warrant for Moore-2 is immediate. That may not be enough to dispel the illusion, but I hope it will begin to counteract it.

Let's confront an important fourth question. What if a subject doesn't *believe* the skeptic's conclusion, but is merely *agnostic* as to whether it's true? What if she doesn't *believe* any undermining hypothesis, but merely *entertains doubts* as to whether they're true?²² Should Moorean reasoning be rationally effective against her agnosticism or doubts? Most likely it *won't* be effective, and it can seem like that's rationally appropriate. Someone who considers Moorean reasoning from such a position *shouldn't* be persuaded by it to overcome her agnosticism or doubts. Can the dogmatist give any account of this?

This is difficult, in part because it's difficult to identify what agnosticism and doubts really *are*. To begin, we should distinguish *agnosticism* from the state of *having no doxastic attitude at all* towards an issue. A subject who's never considered a question isn't yet *agnostic* about it; she has no doxastic attitude at all. Agnosticism is instead a particular kind of attitude, one of suspending belief. As an initial approximation, we may construe

²² See Wright 1985, p. 437: "One the hypothesis is seriously entertained that it is likely as not, for all I know, that there is no material world as ordinarily conceived..." See also his 2007, §IV.2, and White 2006's three-card scenario.

agnosticism as a kind of mixed partial belief: you partly believe things are one way, and partly believe they're the other way. This is only an approximation: as Scott Sturgeon has argued, agnosticism seems to involve a kind of "committed neutrality" that a mere state of mixed belief may lack. Still, it will be enough to start with.

As an initial approximation, "entertaining doubt" may also be construed as having a kind of mixed belief. (Or perhaps it's not mixed *real* belief, but instead mixed *pretend* or *suppositional* belief.) This too is only an approximation, but again I think it will be enough to start with.

Now undermining, like support, is a matter of degree. If a subject who *all-out* believes an undermining hypothesis thereby hypothetically undermines his perceptual warrant, a subject who partially believes (or partially pretend-believes) an undermining hypothesis will thereby hypothetically undermine his perceptual belief *to some degree*.

So the same mechanisms we've been discussing will be operative; and I think they'll still contribute to the illusion that we need antecedent warrant to believe things aren't as the skeptic proposes.

At one point Wright says:

<EXT>

I cannot *rationally* form the belief that it is currently blowing a gale and snowing outside on the basis of my present visual and auditory experience while simultaneously agnostic, let alone skeptical, about the credentials of that experience.²³

I agree with him; there's nothing in this passage that a Moorean dogmatist needs to resist. When you form any belief, it's irrational to *combine* the belief with a continuing agnosticism

²³ Wright 2004, p. 193; compare 2003, p. 61; 2007, §§ II, IV.1, and IV.2; 2000a, p. 142; and 2002, pp. 334–5.

about the belief's warrant. Believing on the basis of some warrant would be hypothetically undermined by such agnosticism. It *doesn't follow* that you need antecedent warrant to believe that your perceptual beliefs are warranted after all (that your experiences have good "credentials"). For all Wright says here, subjects who have *no doxastic attitudes whatsoever* about the credentials of their experience, or subjects who just *uncritically accept that* their experiences warrant their perceptual beliefs, may be reasonable in holding their perceptual beliefs. Not because that acceptance is covertly warranted—but rather because they merely *lack* the doxastic attitudes that can hypothetically undermine the warrant their experiences are immediately providing. Being perceptually warranted doesn't require us to *have* antecedently *warranted positive* higher-order beliefs, to the effect that our experiences *do* give us warrant. But as we're seeing, it may require us to *lack negative* higher-order attitudes—warranted or not—that hypothetically undermine, and so rationally interfere with, our first-order, immediate perceptual warrants.²⁴

I don't expect these explanations to satisfy everyone. In particular, if you're sure it must be *less reasonable* for the subject to believe Moore-3 on the basis of Moore-2 than it is for her to believe Moore-2, these explanations probably won't satisfy you. So be it. They at any rate constitute *some* accounting for why Moorean reasoning should seem sketchy to us, compatible with its retaining its capacity to transmit warrant. I don't expect what I've said here to be decisive.

²⁴ This bears on the "Simple Elevation Hypothesis" Wright formulates in his 2007, §II. Roughly, that principle says that if *C* are the conditions that confer some warrant *W*, then being in a position to warrantably believe *C* obtains suffices for being in a position to warrantably believe you have *W*. I agree that *warrant to believe C obtains may warrant* some subjects in believing they have *W*. But it can't *suffice to put them in a position to warrantably believe* they have *W*. They'd need also to *lack* interfering higher-order attitudes, including agnosticism.

<H1>VI

How does all this bear on Wright's own understanding of transmission-failure?

Wright—and Davies also—alternate between different glosses on what transmission-failure consists in. One criterion they give is: Can reasoning through an argument, relying on the indicated warrants for its premises, add up to (or constitute, or give you) *any* new warrant to believe the argument's conclusion?²⁵ Another family of criteria asks things like: Can reasoning in the indicated way *rationally resolve doubt or agnosticism* about the argument's conclusion?²⁶

In Section V, above, I reviewed reasons for thinking that Moorean reasoning might count as transmission-failing by the second criterion even if it didn't by the first—that is, even if it *did* give subjects new warrant for Moore-3, and the right kinds of subjects (subjects without interfering higher-order attitudes) would be able to reasonably *believe* Moore-3 on the basis of such reasoning. So if the notion of transmission-failure is to track an argument's inability to do that, then we should take the first criterion as definitive, and set the second aside.²⁷

But when does that first criterion apply? What sorts of arguments can and what sorts

²⁵ See Wright 2000a, pp. 142–3, 154; 2002, pp. 335, 342; 2007, §III; Davies 2000, pp. 388, 393, 394, 399, 400. Sometimes this is put differently: Can reasoning in the indicated way add up to a *first* warrant to believe the argument's conclusion? (Wright 2000a, pp. 141, 155; 2002, pp. 331–2; 2003, p. 57; Davies 2000, pp. 386, 388). It's unclear whether every argument that passes the main criterion for transmitting warrant needs also to pass that variant (see Davies 2000, p. 388 and fn. 17).

²⁶ Wright 2000a, p. 140; 2002, p. 331; Davies 1998, pp. 348–9; Davies 2000, pp. 386, 387, 388, 394, 396–7, 398–9, 400, 401. A significantly different variant is: Can doubt about the conclusion *be rationally combined with* preparedness to reason in the indicated way? See Davies 2000, pp. 397, 398–9, 401, 402–4, 410.

²⁷ Though see Davies's development of the second criterion in Davies 2009.

can't transmit warrant?

Wright has addressed that question in several places. In most of his discussion he seems to be operating with some form of our BW Model of transmission-failure. He does distinguish between two transmission-failure “templates.” He calls them the *information-dependent* and the *disjunctive* template. (I'll indicate Wright's terminology with this typeface.) These present some interpretive difficulties, but as we'll discuss, they might properly be understood as two subspecies of the BW Model. So it's not clear that Wright countenances any forms of transmission-failure other than the BW Model. If he doesn't, then he should agree that if dogmatism is true (and we set probabilistic worries aside) Moorean reasoning should be allowed to transmit warrant.

However, there are also elements in Wright's discussion that do suggest non-BW models of transmission-failure.

Throughout his discussion, Wright deploys a contrast between *information-dependent* and *non-inferential* warrants. Wright says that what makes a warrant *non-inferential* is that the warranted belief not be *based on any inference*, and that, furthermore, it not draw any of its support from “evidence” or “collateral information.”²⁸ In his earliest discussions, he uses “evidence” to mean “empirical evidence” (see 2000a, p. 147). In later discussions, he counts earned a priori warrants as evidence too (2004, pp. 174–5). But nowhere does he count unearned warrants as evidence. So now consider beliefs whose warrant is non-inferential, in the sense I articulated in section III, but still mediate, because it depends on antecedent *unearned* warrant to believe other things. Wright will regard these beliefs as *non-inferential* and “groundless” (2000a, pp. 152–3, 156–7; and 2003, pp. 66–67). In his view, that includes our most basic perceptual beliefs. In this table:

²⁸ For his definitions of “information-dependent,” see Wright 2000a, p. 143; 2002, pp. 335–6; and 2003, p. 59. On “non-inferential,” see Wright 2002, p. 341 and 2003, pp. 60, 62.

What I call immediate warrants	What I call mediate but non-inferential warrants	What I call inferential warrants
	Non-inferential warrants that depend on antecedent <i>unearned</i> warrants to believe other things	
	Non-inferential warrants that depend on “information” (that is, antecedent <i>earned</i> warrants to believe other things)	

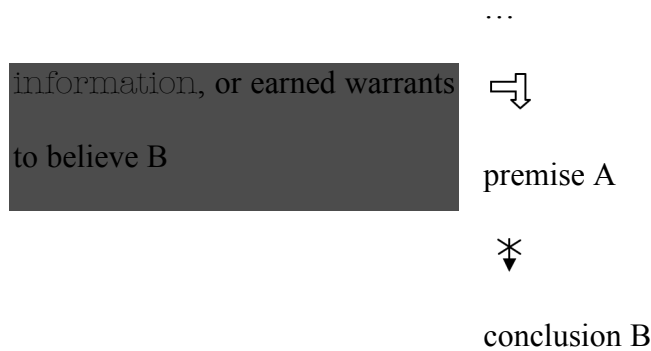
the darkest gray cells represent Wright’s information-dependent warrants; the other cells his non-inferential warrants.

In Wright’s vocabulary, our perceptual warrants have a Humean or “I-II-III” structure when they’re “inferentially” supported by “defeasible evidence” (see 2002, pp. 336 and 341; 2003, p. 60; 2004, p. 172). They have only a Cartesian structure if they’re “acquired not on the basis of evidence, but directly” (2003, pp. 60, 62; 2004, p. 168). This bears on what sorts of skeptical argument, and what sorts of transmission-failure, our perceptual beliefs are vulnerable to.

As I said, Wright outlines two templates for transmission-failure. The first is the information-dependent template. This is instantiated when some premise in a piece of reasoning is warranted in part by information, that is, antecedent earned warrants, that


includes the reasoning's conclusion. That's clearly an example of our BW Model of transmission-failure:

<DIS>



But the BW Model also would apply to cases where your reasoning depends on antecedent *unearned* warrants that include the reasoning's conclusion:

<DIS>



An important question is whether this is all that Wright's second template, his *disjunctive* template, manages to cover.²⁹ If so, then Wright's two transmission-failure templates are just subspecies of our BW Model.

However, Wright *seems to* want his *disjunctive* template to cover cases where the

²⁹ For Wright's exposition of the *disjunctive* template, see Wright 2000a, pp. 153–6; 2002, pp. 342–4; and 2003, pp. 60–3.

relevant warrants are alleged to be genuinely immediate. So perhaps the disjunctive template only overlaps with our BW Model, and includes more besides.

It's not easy to determine which of these two possibilities obtains. My considered view is that the former is correct: that the only models of transmission-failure Wright successfully presents us with are BW Models. When he's trying to apply his disjunctive template to cases where your warrant for the relevant premises is *alleged to be* genuinely immediate, he's most charitably interpreted as arguing that the warrant in question *can't* really be immediate, but can at best be of the sort represented by the light gray cell in our table. (I'm not sure whether this captures Wright's own understanding of what he's doing.)

I'll briefly present some evidence in support of this interpretation. But as I said, the question is tricky. (I'll reserve some of the harder evidence until later.)

To begin, Wright at some points (in 2003, at note 9 and pp. 67–8) says that transmission-failure *presupposes* Closure. That is, in order for us to have a case where you have warrant for A that fails to transmit to B, it must be that Closure is not violated: you must also have some warrant to believe B. Why? Wright seems to be thinking that if your warrant for A didn't already *require* that warrant for B, you wouldn't be in the kind of position that blocked warrant-transmission. In other words, transmission-failure must be due to B's being among the background warrants that warrant you in believing A. This is just our BW Model.

Let's next consider Wright's expositions of his disjunctive template. Wright is supposing that he has some warrant W his possession of which is compatible with each of A and C being the case. He believes A on the basis of that warrant, and reasons from A to B. B is such that its falsity implies C is true; to simplify we can just set B equal to not-C. Wright focuses on the question: what should make it reasonable for him to believe A on the basis of W, rather than the more tentative belief $A \vee C$? Here is a sample of what he says:

<EXT>

[T]he more tentative claim would indeed be appropriate unless I am somehow *additionally* entitled to discount alternative C . . . [I]n order for me to be entitled to discount C, and so move past the disjunction to A, I have to be entitled to *discount the negation of B*, and therefore entitled to accept B; for by hypothesis, if not-B were true, so would C be. So it would seem that I must have an appreciable entitlement to affirm B *already*, independent of the recognition of its entailment by A, if I am to claim to be warranted in accepting A in the first place . . . So my warrant [for A] does not transmit [to B].^{30 31}

Notice what he says: for the subject to reasonably believe A on the basis of W, he needs to be antecedently “entitled to discount” not-B, and “entitled to accept” and “affirm” B. Elsewhere

³⁰ Wright 2003, pp. 62–3, Wright’s italics; compare his 2000a, p. 155 and 2002, p. 343).

³¹ An aside: Wright says in a note that this argument relies on a closure step: if you need antecedent warrant to believe not-C (because your warrant for A draws support from that antecedent warrant), and you appreciate that not-B implies C, then you *will also need* some warrant to believe B. However, I think that’s not enough for Wright’s purposes. The fact that you must in this situation have *some* warrant to believe B does not entail that this warrant has to be antecedent to A, nor that your warrant for A will constitutively depend on it too. So we wouldn’t yet have reason to deny that warrant could be transmitting from A to B. Wright is on better footing if he assumes that your antecedent warrant for not-C not merely *implies you have* some warrant for B; but assumes, more strongly, that your warrant for not-C *will transmit* to B. Then it will follow, in the situation he describes, that you have *a* warrant for B which is antecedent to A. I think even more will be needed to clinch his case: having *a* not-C-derived, antecedent-to-A warrant for B is compatible with *also* having an A-derived warrant for B. So Wright ultimately needs to show that the warrant you have for A *includes* the not-C-derived, antecedent to-A warrant for B. That would make it clear why your warrant for A should be blocked from transmitting to B. This will be easiest when B implies, or otherwise supports, not-C. When B is merely implied by not-C, more work will be needed.

Wright talks of having “antecedent reason to think” that B holds (2000a, p. 156; see also 2003, p. 66 and 2002, p. 346). In other words, the subject’s warrant for A must require him to have an antecedent warrant to believe or doxastically accept B. Presumably, then, the warrant for A is not immediate. But it can still be *non-inferential* in his sense: it can be a warrant from **the light gray cell** in our table.

That’s the most straightforward construal of Wright’s *disjunctive template*, and if it’s correct, then this template will be just a subspecies of our BW Model. It won’t apply in cases where your warrant for A genuinely *is* immediate:

<DIS>

[experiences]



A



B

However, there are reasons to hesitate: notice in the second-to-last sentence of the quoted passage, Wright doesn’t say he needs an antecedent warrant to believe B *if he is to be* warranted in accepting A. He says he needs it in order to *claim to be* warranted in accepting A. Is this difference significant? Should we take Wright to be talking not about what it takes to be warranted in believing A, but rather about what it takes to have some *higher-order warrant*, to the effect *that one is* warranted in believing A?

If the *claiming warrant* talk was supposed to mean that, Wright didn’t do much to highlight it. He mostly writes as though being in a position to *claim a warrant* were just a matter of *having* a warrant—or perhaps, having an *earned* warrant. For example, he says:

<EXT>

Whenever I get in a position to claim justification for a proposition, I do so courtesy of specific presuppositions . . . for which I will have no specific, earned evidence . . . Again: whenever cognitive achievement takes place, it does so in a context of *specific* presuppositions which are not themselves an expression of any cognitive achievement to date.³²

He speaks here as though getting in a position to claim justification were the same thing as “cognitive achievement tak[ing] place.” Yet he uses “cognitive achievement” to cover getting evidence for first-order claims like P, or coming to know P (see e.g., 2004, p. 175). It’s not said to require *higher-order* warrants.

Wright elaborates on how he understands the notion of claiming warrant in his 2007. There he confirms that it should not be understood as *just equivalent to* believing, or acquiring warrant to believe, that one has a warrant. Instead, the main role of the notion is to express a kind of *internalist* warrant. It’s unclear to me whether he wants being in a position to claim warrant to at least always *include* having a higher-order warrant.

There are other higher-order threads in Wright’s discussion. Consider these corresponding sentences from his different expositions of the disjunctive template:

<EXT>

[W]hat, in the circumstances, can justify me in accepting A?

Why not just reserve judgment and stay with the more tentative disjunction, “A or C”?³³

<EXT>

³² Wright 2004, p. 189.

³³ Wright 2000a, p. 155.

[W]hat, in the circumstances, can justify me in accepting A?
Should I not just reserve judgment and stay with the more
tentative disjunction, either (I have warrant for) A or C?³⁴

<EXT>

[W]hat, in the circumstances, can justify me in taking it that I
have a warrant for A? Why not just reserve judgment and stay
with the more tentative disjunction, Either (I have a warrant
for) A, or C?³⁵

Wright doesn't say anything to highlight these changes; and neither does he say anything to *motivate* focusing on higher-order questions in this context. Moreover, his dialectical settings demand that we *shouldn't* be focusing on higher-order questions. Wright's expositions begin with questions about whether you've acquired warrant to believe that *certain animals are not disguised mules*, or that *the wall you see is not white but lit by tricky red lighting*. He says that you may warrantably be taking it for granted that the italicized claims are true, but it would be "absurd to pretend that you had gained a reason for thinking so." He then says his *disjunctive template* is suggested by, or is a generalization of, what's going on in such examples. Up to that point, everything was formulated in terms of what *first-order* warrants we possess.

So I don't think there's enough in the texts to support reading much into Wright's higher-order-ish talk. But we'll revisit this issue in the next section.

We've been discussing how to interpret Wright's notion of *non-inferential* warrant, and his *disjunctive* transmission-failure template. There are also some other, more minor,

³⁴ Wright 2003, 62; my underlining.

³⁵ Wright 2002, 343; my underlining.

interpretive issues to sort out.

One of these is that, in his 2004, Wright hesitates to speak of being entitled *to believe* anything. What's driving him is a concern whether the notion of "belief" can be used in a way that permits one to rationally believe things on mere "entitlements," that is, unearned warrants (2004, pp. 175–7). When the only warrants one has are unearned, perhaps we should instead speak of *trusting* that things are as those warrants testify. If so, then Wright's account of perception should be revised to say: our perceptual warrants require us to have antecedent entitlements to trust that our senses are reliable, that we're not brains in vats, and so on.

This talk of entitlements to trust might evoke the notion of "taking for granted" that we discussed in §III, above. But that's not what Wright means. His notion of *trusting* that P is still a species of "accepting," which is a general kind of action-guiding attitude. *Trusting* is like belief in that it's a form of acceptance that excludes doubt and withholding judgment (2004, pp. 192–4). At one place Wright characterizes entitlement as a "norm of doxastic acceptance" (2004, note 1). So I think we can group belief and trust together as two species of doxastic acceptance. The sense in which a *dogmatist* says you "take it for granted" that you're not a brain in a vat, and so on, is instead just being vulnerable to that as a defeater. It's no kind of doxastic attitude towards the hypothesis that you're not a brain in a vat.

If Wright's notion of *trusting* is still a form of doxastic acceptance, then I don't think the contrast between it and belief bears much on the issues we're discussing. If there are differences between different kinds of doxastic acceptance, they're not differences I care about. I'm interested in the general question of what we have warrant to doxastically accept in any form.

Another interpretive issue concerns the fact that Wright seems to understand entitlements as warrants that don't "bear on the likely truth" of what they support (2004, pp.

177, 182). This is puzzling. Most epistemologists would agree that there are norms of doxastic acceptance that don't bear on the likely truth of one's belief, but they'd deny that these are properly regarded as *epistemic warrants*. (That doesn't have to mean they have no relevance to epistemology.) For instance, Pascalian considerations might mandate believing or at least doxastically accepting that God exists, but they wouldn't do anything to make that hypothesis epistemically more probable, and so they shouldn't be counted as *epistemic warrants* that God exists. I've been assuming that Wright does intend his entitlements to be a kind of epistemic warrant. But I'm not entirely sure that's correct. I'd find the indispensability considerations he adduces in his 2004 more convincing if they *weren't* supposed to constitute epistemic warrants (see also 2000a, p. 157; 2003, pp. 66–7, 68). Perhaps this is a sign that Wright and I are operating with fundamentally different understandings of what an epistemic warrant needs to be. Well, I'll hope there's enough overlap in our understanding that I haven't seriously misrepresented him.

Summarizing: based on the evidence we've reviewed so far, Wright's two transmission-failure templates seem to be subspecies of our BW Model. As we've noted many times, that model can't apply when your warrant for the relevant premises is genuinely immediate.

<H1>VII.

Could there be other models of transmission-failure, which apply even when your warrant for the premises is immediate?

Our Background Warrant Model says:

<DIS>

W

... B \Rightarrow *if B plays this role,*
 A
 \nrightarrow *then this reasoning doesn't transmit warrant*
 B

If B is among the background you need antecedent warrant for, *to be warranted by W to believe A*, then your W-based warrant for A can't transmit to B.

Some of Wright's higher-order-ish talk suggests a different, **Higher-Order Background Warrant (HOBW) Model:**

<DIS>

	<i>if B comes here ...</i>	W
... or here	\Rightarrow	\Downarrow
	W warrants believing A	A
		\nrightarrow <i>then this reasoning doesn't transmit warrant</i>
		B

If B is among the background you need antecedent warrant for, *to be warranted in believing you are warranted by W to believe A*, then your W-based warrant for A can't transmit to B.

The suggestions are strongest in Wright's 2002 exposition of his disjunctive template. Later

in that same paper he also writes:

<EXT>

[A] warrant, W, for a belief, A, cannot transmit to any of its consequences, B, if—in context—one would need an entitlement (earned or standing) to B in order to defend the claim that conditions for the acquisition of W were satisfied.³⁶

In other words, transmission from A to B fails when an antecedent warrant for B is needed to support the higher-order claim that you have acquired (or are at least in a position to acquire) warrant W to believe A.

What could motivate such a restriction on warrant-transmission? After the work of Alston and Audi,³⁷ there's now broad agreement among epistemologists that higher-order epistemic statuses may be harder to achieve, and may need to draw on more resources, than their first-order counterparts. So why should facts about what's needed to achieve the higher-order warrant that *you are warranted by W to believe A* bear on what the first-order warrant for A is itself able to transmit to?

Even if we were to *grant* the HOBW Model, it's not clear to me that Moorean reasoning as understood by dogmatists would instantiate it. Some dogmatists, like myself, are internalists about warrant. They should say: look, for W to warrant you in believing you have hands, it's not a requirement that your senses really *be* reliable, or that you really *fail* to be a brain in a vat. So why, for *you to warrantedly believe that W so warrants you*, should it be a requirement that you have antecedent warrants to believe those things? It ought to be enough that you merely *lack* warrants to believe their negations. More needs to be done to persuade

³⁶ Wright 2002, p. 345.

³⁷ See the essays in Alston 1989, especially Essay 6; and Audi 1993, especially Chapters 4 and 11.

us internalist dogmatists that antecedent warrant for Moore-3 is *even needed for the higher-order claim* that your experiences warrant you in believing Moore-2.

What may make the HOBW Model attractive is the hovering thought that dogmatism can at best be an answer to first-order skepticism; that it's going to let second--order skeptics have run of the house.³⁸ This thought goes: even if a lack of perception-independent warrant to believe we're not brains in vats doesn't preclude us from *having* perceptual warrants, it *will* preclude us from *warrantedly believing* we have them. Dogmatists can at best give us "animal warrants" to believe we have hands; not warrants that survive reflective scrutiny.

But I see no reason to concede that. Why can't our immediate perceptual warrants be available to us for settling higher-order questions, too?

What *is* true is that it's easier to give an accounting of our first-order perceptual warrants than it is to give an accounting of our higher-order warrants. Our first-order warrants are constituted by everyday experiences. Our higher-order warrants are constituted by lengthy philosophical discussions, of which the present discussion is just one piece. But if dogmatism is correct, shouldn't it be possible for us (eventually) to be warranted in believing it is? (These philosophical discussions ought to lead somewhere!) Why can't our immediate perceptual warrants be part of our (more dearly won) warrant to believe dogmatism?

I suppressed the HOBW Model of transmission-failure until now for expository convenience. It was useful for us to get clear on the prospects of the BW Model first. But as I've just argued, I don't think the HOBW Model necessarily changes our dialectical position.

What I think will be more promising to consider are some other models, which try to motivate obstacles to Moorean transmission *without* making claims about what higher-order warrants we're in a position to have. These other models may, on some views, end up *entailing* the HOBW Model. But since they don't *rest on* any contested claims about what

³⁸ See Wright 1985, p. 444; 1991, p. 89; and 2007, §II.

higher-order epistemic positions we're in, they may be dialectically more useful.

In the framework we're working with, warrants are constituted by factors that count in favor of believing something. Let's call these **warrant-making factors** (on analogy with the notion of "truth-makers"). Warrant-making factors may include things like: having antecedent warrant to believe such-and-such premises. They may include things like: having such-and-such experiences. On reliabilist theories, they will include things like: having reliable vision. And so on.

This enables us to formulate one new model of transmission-failure.

<EXT>

Warrant-Making Factor (WMF) Model: If B is among the factors that *make* W be a warrant for A, then your W-based warrant for A can't transmit through reasoning to the conclusion that B obtains.³⁹

This is different than the BW and HOBW Models, because B need not include your having warrant to believe anything. B may instead be a fact like: your vision is reliable, or: you have such-and-such experiences. For example, the WMF Model would count this reasoning as transmission-failing:

<DIS>

[Visual experiences as of hands]



³⁹ Compare Wright's repeated insistence that warrant cannot transmit to the "external preconditions for the effectiveness of your [belief-forming] method" (2000a, p. 154; 2002, p. 342; 2003, pp. 61, 70; see also his remarks on "operational necessity" in the same papers). See also Davies 1998, p. 354 and 2000, p. 404; and Principle QB in McLaughlin 2000.

I see two hands.



I have visual experiences.

In the Gettier literature, one sees the term “defeat” used in such a way that defeating facts merely have to *obtain*, in order to neutralize an epistemic status you’d otherwise possess. I use the term differently. I only use “defeat” to describe a kind of *warrant* or evidence subjects get. When warrant-makers give you prima facie warrant to believe A, that warrant is vulnerable to being outweighed or undermined by other **defeating evidence**. Let’s focus on the undermining kind of defeat. This is exemplified in the following sort of case: you seem to see Tom at the library. Then you learn that Tom is one of a set of identical triplets. Intuitively, this new information doesn’t tell you that Tom *wasn’t* at the library. At least, not directly. It might on balance be reasonable for you to reduce your confidence that Tom was there. But what the new information more directly testifies to is that *you weren’t in a position to visually tell* that Tom was at the library. In doing that, it *undermines* your initial visual warrant to believe Tom was there.

When a hypothesis U is such that *evidence for it* will undermine a warrant W for A, I call U an undermining hypothesis for W. Such hypotheses lower the likelihood that in having W you’d thereby be *perceiving*, or more generally, in a *position to tell*, that A. On the other hand, I’ll call a hypothesis an **anti-underminer** with respect to W and A when (i) it entails that some underminers of W’s support for A are false; and (ii) it’s not itself such an underminer.

This enables us to formulate another new model of transmission-failure:

<EXT>

Anti-Underminer (AU) Model: If B is an anti-underminer

with respect to W and A, then your W-based warrant for A can't transmit through reasoning to the conclusion that B.⁴⁰

Things can be said to motivate each of these new models. As we saw in §III, above, it can be natural even for a dogmatist to describe you as “taking for granted” that your senses are reliable, and so on. And it can be natural to worry whether arguments can be effective whose conclusions are already “taken for granted,” or “presupposed,” in that way.⁴¹

Though these proposals can be motivated, I do not think they can be correct. I have counter-examples.

Consider:

<DIS Proposition>

Turing-1. Being able to grasp entailments is sufficient for being able to rationally follow an argument.

Turing-2. Turing machines are able to grasp entailments.

Turing-3. So Turing machines are able to rationally follow arguments.

⁴⁰ See Wright's characterization of his “triads” in 2007.

⁴¹ Wright usually uses the vocabulary of “presuppositions” to refer to antecedently warranted background premises: see for example 2000a, p. 148; and his repeated references to presupposing *warrants* and *entitlements*: 2000a, p. 159; 2002, p. 345; 2003, pp. 64, 66, 75. See also Davies 1998, pp. 350–2, 354. However, as Wright points out in 2007, §I, “presupposition” can also be understood in the more lightweight way being envisaged here. See also the references in fn. 9, above; and Alston 1986, who says that when I form perceptual beliefs, I'm “assuming” or “proceeding as if” sense perception is reliable (pp. 323–4, 327–9), without needing to have any antecedent warrant to believe that it is (pp. 331–2).

Never mind how we might establish Turing-2. Just assess whether this argument seems to be objectionably question-begging. Most of my consultants judge it is not.

And yet, if you learned that *you hadn't properly followed* the argument, that would undermine how confident you ought to be in its conclusion (at least, on the basis of so reasoning). If you learned that you *weren't capable* of following the argument, all the more so. Now it's certainly epistemically possible that you are merely a Turing machine, so if you acquired reason to believe (let us suppose, falsely) that mere Turing machines are never able to rationally follow arguments, it should to some degree have that undermining effect. What that means is that the conclusion of this argument is one kind of anti-underminer for the warrant the argument gives you. (It's an anti-underminer for the warrant you get by *reasoning through* or following the argument, not your warrant for any of the premises.) So the AU Model would count this argument as transmission failing. Yet intuitively, it seems like such reasoning should be okay.

If in fact *you are* merely a Turing machine, then the conclusion of the argument may be not just an anti-underminer, but a warrant-making factor as well. In that case, the WMF Model would also count the argument as transmission-failing. But again, it seems like it ought not to be so counted.

Consider a different case. You're looking in a zoo pen, and your mad scientist friend proudly announces that any animals in the pen have been turned invisible. You inform him he's wrong: you see rats in the corner, so there are some visible animals in the pen after all. Your reasoning might be represented like this:

<DIS Proposition>

Rat-1. [Visible experiences as of rats]

Rat-2. I see rats in the corner.

Rat-3. So there are some visible animals in the pen.

Assess whether this argument seems objectionably question-begging. Most of my consultants judge it is not.

And yet, to the extent that you have reason to believe any animals in the pen *are* invisible, your visual warrant to believe Rat-2 is undermined. So Rat-3 is an anti-underminer for the warrant Rat-1 gives you to believe Rat-2. So the AU Model would count this argument as transmission-failing, too. Yet intuitively, it seems like such reasoning should be okay.

The natural move at this point will be to try to refine the transmission-failure models to avoid these counter-examples. For example, we might try to replace the AU Model with:

<EXT>

If B is an anti-underminer for W *whose non-obtaining would give you the same experiences you actually have* then your W-based warrant for A can't transmit through reasoning to the conclusion that B.

This would no longer count the Rat argument as transmission-failing, since if the pen *did* only contain invisible animals, we wouldn't expect to have experiences of rats. It would only impugn arguments like this:

<DIS Proposition>

Rat-1. [Visible experiences as of rats]

Rat-2. I see rats in the corner.

Rat-3*. So it's not the case that: there are no visible animals in the pen and I am merely hallucinating rats.

and that seems close enough to Moorean reasoning that we have no independent intuition it should be okay.

Perhaps some such refinement can work. But I'm pessimistic, for three reasons.

In the first place, the refinement I articulated wouldn't help with the Turing example.

In the second place, what may make such refinements attractive are the probabilistic considerations that I deferred for discussion elsewhere. Those probabilistic considerations allegedly threaten the possibility of something's warranting you in believing Q, when not-Q's truth would entail that you have the same warrant. I acknowledge the prima facie force of those considerations; and I have responses to them; but they require extended discussion. Our question here is whether the notion of transmission-failure gives us any further diagnostic tool, or whether it's just piggy-backing on the probabilistic considerations. I suspect that efforts to refine the AU Model in the way I described will lead to our just piggy-backing.

As I said in section IV, I doubt that the probabilistic considerations really do reflect our deepest-rooted intuitive resistance to Moorean reasoning, anyway, for we feel that same resistance to arguments where the probabilistic considerations are not applicable.

My third doubt about the strategy of refining the AU Model is that it raises serious motivational difficulties.

It was possible to motivate the AU and WMF Models in the first place because they reflected notions of "taking for granted" or "presupposing" that we already had some intuitive grasp of: it can seem natural to doubt whether arguments can establish their own "presuppositions." That can seem like a covert form of question-begging.

This motivation wouldn't have been available if we didn't have some folk understanding of anti-underminers or warrant-making factors *as* "presuppositions," and had just used some technical notion instead. It *doesn't* seem natural to doubt whether arguments can establish their own XYZZYs. The refined AU Model strikes me as going down that path.

I don't think we have any intuitive notion of "presupposing" such that the relation specified in the original AU Model and the relation specified in the refined model count as *different kinds of presupposing*. No, the sense in which one can naturally be counted as "presupposing" B in the refined model is *the same* as in the original model; it's just that in the refined model, B meets extra conditions besides. Yet our counter-examples show that it's *a mistake* to count "presupposing" your conclusion in that sense as any form of begging a question. The proponent of the refined model needs to scratch up some further motivating story. I'm not sure there is any (non-probabilistic) such story.

Perhaps all that will be said for a refined story is that it extensionally captures the kinds of arguments, like Moore's, that one doesn't like.

I understand and respect the view that, when B is the kind of conclusion we're considering, the subject's warrant W for A *can't* be immediate, but instead requires her to have antecedent warrant to believe B. As a dogmatist, I don't *agree* with that view. But I recognize it as a viable theoretical option. In that case, Moorean reasoning would fail because it exemplifies the simple BW Model.

If on the other hand we evaluate Moorean reasoning from the dogmatist's perspective, it seems to me more promising to diagnose our intuitive resistance to the reasoning as the product of illusion. I don't see much money in the project of coming up with (non-probabilistic) models of transmission-failure that impugn Moorean reasoning even when our perceptual warrant *is* immediate.

<REF>Bibliography

Alston, W. 1989 *Epistemic Justification*, Ithaca, Cornell University Press

Audi, R. 1993 *The Structure of Justification*, Cambridge, Cambridge University Press.

- Bergmann, M. 2004 “Epistemic Circularity: Malignant and Benign”, *Philosophy and Phenomenological Research* 69, pp. 709–27.
- Bergmann, M. 2005 “Defeaters and Higher-Level Requirements”, *The Philosophical Quarterly* 55, pp. 419–36.
- Broome, J. 1999 “Normative Requirements” *Ratio* 12, pp. 398–419.
- Burge, T. 1993 “Content Preservation”, *Philosophical Review* 102, pp. 457–88.
- Burge, T. 1996 “Our Entitlement to Self-Knowledge”, *Proceedings of the Aristotelian Society* XCVI, pp. 91–116.
- Burge, T. 2003 “Perceptual Entitlement”, *Philosophy and Phenomenological Research* 67, pp. 503–48.
- Cohen, S. 1988 “How To Be a Fallibilist”, *Philosophical Perspectives* 2, pp. 91–123.
- Cohen, S. 1999 “Contextualism, Skepticism, and the Structure of Reasons”, *Philosophical Perspectives* 13, pp. 57–89.
- Cohen, S. 2000 “Contextualism and Skepticism”, *Philosophical Issues* 10, pp. 94–107.
- Conee, E. and Feldman, R. 2004 *Evidentialism*, New York, Clarendon Press.
- Davies, M. 1998 “Externalism, Architecturalism, and Epistemic Warrant”, in C. Wright, B. C. Smith, and C. Macdonald (eds.) *Knowing Our Own Minds*, Oxford, Oxford University Press, pp. 123–54.
- Davies, M. 2000 “Externalism and Armchair Knowledge”, in P. Boghossian and C. Peacocke (eds.) *New Essays on the A Priori*, Oxford, Oxford University Press, pp. 384–414.
- Davies, M. 2003a “The Problem of Armchair Knowledge”, in S. Nuccetelli (ed.) *New Essays on Semantic Externalism and Self-Knowledge*, Cambridge, Mass., MIT Press, pp. 23–55.
- Davies, M. 2003b “Externalism, Self-Knowledge, and Transmission of Warrant”, in M. Frápolli and E. Romero (eds.) *Meaning, Basic Self-Knowledge, and Mind: Essays on Tyler Burge*. Stanford, CSLI Publications, pp. 105–30.

- Davies, M. 2004 “Epistemic Entitlement, Warrant Transmission, and Easy Knowledge”, *Proceedings of the Aristotelian Society* Supp. Vol. LXXVIII, pp. 213–45.
- Davies, M. 2009 “Two Purposes of Arguing and Two Epistemic Projects”, in Ian Ravenscroft (ed.) *Mind, Ethics, and Conditionals: Themes from the Philosophy of Frank Jackson*, Oxford, Oxford University Press, pp. 337–383.
- Dretske, F. 1970 “Epistemic Operators”, *Journal of Philosophy* 67, pp. 1007–23.
- Dretske, F. 2000 “Entitlement: Epistemic Rights without Epistemic Duties?” *Philosophy and Phenomenological Research* 60, pp. 591–606.
- Field, H. 2000 “A Prioricity as an Evaluative Notion”, in P. Boghossian and C. Peacocke (eds.) *New Essays on the A Priori*, Oxford, Oxford University Press.
- Field, H. 2005 “Recent Debates about the A Priori”, *Oxford Studies in Epistemology* 1, pp. 69–88.
- Ginet, C. 1980 “Knowing Less By Knowing More”, *Midwest Studies in Philosophy* 5, pp. 151–61.
- Hawthorne, J. 2004 *Knowledge and Lotteries*, Oxford, Oxford University Press.
- Klein, P. 1981 *Certainty: a Refutation of Skepticism*, Minneapolis, University of Minnesota Press.
- McKinsey, M. 2003 “Transmission of Warrant and Closure of Apriority”, in S. Nuccetelli (ed.) *New Essays in Semantic Externalism and Self-Knowledge*, Cambridge, Mass., MIT Press, pp. 97–116.
- McLaughlin, B. 2000 “Self-Knowledge, Externalism, and Skepticism”, *Aristotelian Society Supplemental Volume* 74, pp. 93–117.
- Meiland, J. 1980 “What Ought We to Believe? The Ethics of Belief Reconsidered”, *American Philosophical Quarterly* 17, pp. 15–24.
- Peacocke, C. 2004 *The Realm of Reason*, New York, Oxford University Press.

- Plantinga, A. 1993a *Warrant: The Current Debate*, New York, Oxford University Press.
- Plantinga, A. 1993b *Warrant and Proper Function*, New York, Oxford University Press, New York.
- Pollock, J. 1983 “Epistemology and Probability”, *Synthese* 55, pp. 231–52.
- Pryor, J. 2000 “The Skeptic and the Dogmatist”, *Noûs* 34, pp. 517–49.
- Pryor, J. 2004 “What’s Wrong With Moore’s Argument?”, *Philosophical Issues* 14, pp. 349–78.
- Pryor, J. 2006a “Hyper-Reliability and A Priority”, *Proceedings of the Aristotelian Society* CVI, pp. 327–44.
- Pryor, J. 2006b “More on Hyper-Reliability and A Priority”, available online at <<http://www.jimpryor.net/research/papers/More-Hyper.pdf>>.
- Pryor, J. (ms-a) “Problems for Credulism”, available online at <<http://www.jimpryor.net/research/papers/Credulism.pdf>>.
- Pryor, J. (ms-b) “Hypothetical Oughts”, available online at <<http://www.jimpryor.net/research/papers/Hypothetical.pdf>>.
- Silins, N. 2005 “Transmission Failure Failure”, *Philosophical Studies* 126, pp. 71–102.
- Sturgeon, S. 2012 “Confidence and Coarse-Grained Attitudes” *Oxford Studies in Epistemology* 4.
- White, R. 2006 “Problems for Dogmatism”, *Philosophical Studies* 131, pp. 525–57.
- Williamson, T. 2000 *Knowledge and Its Limits*, Oxford, Oxford University Press.
- Wright, C. 1985 “Facts and Certainty”, *Proceedings of the British Academy* 71, pp. 429–72.
- Wright, C. 1991 “Skepticism and Dreaming: Imploding the Demon”, *Mind* 100, pp. 87–115.
- Wright, C. 2000a “Cogency and Question-Begging: Some Reflections on McKinsey’s Paradox and Putnam’s Proof”, *Philosophical Issues* 10, pp. 140–63.
- Wright, C. 2000b “Replies”, *Philosophical Issues* 10, pp. 201–19.

Wright, C. 2002 “(Anti-)Sceptics Simple and Subtle: G. E. Moore and John McDowell”, *Philosophy and Phenomenological Research* 65, pp. 330–48.

Wright, C. 2003 “Some Reflections on the Acquisition of Warrant by Inference”, in S. Nuccetelli (ed.) *New Essays on Semantic Externalism and Self-Knowledge*, Cambridge, Mass., MIT Press, pp. 57–77.

Wright, C. 2004 “Warrant for Nothing (and Foundations for Free)?”, *Proceedings of the Aristotelian Society* Supp. Vol. LXXVIII, pp. 167–212.

Wright, C. 2007 “The Perils of Dogmatism”, in S. Nuccetelli and H. Seay (eds.) *Themes from G. E. Moore: New Essays in Epistemology and Ethics*, Oxford, Oxford University Press.

Zalabardo, J. 2005 “Externalism, Skepticism, and the Problem of Easy Knowledge”, *Philosophical Review* 114, pp. 31-61.