

Problems for Credulism*

Jim Pryor

29 Aug 2012

1. Introduction

I've defended views about the epistemology of perception that I call "dogmatist." They are close to views defended by Michael Huemer and John Pollock, and somewhat more broadly to a range of views defended by others. These views have been criticized for being incompatible with Bayesianism. I will take Roger White's "Problems for Dogmatism" (White 2006) as representing these criticisms.¹ This essay will review and assess some central pieces of the criticisms. Doing so will reveal some limits of Bayesianism when it comes to representing undermining evidence.

I won't argue directly that the criticisms fail, nor concede directly that they are correct. As you'll see, the dialectic is not that straightforward.

Additionally, my aim here isn't to speak on behalf of my own idiosyncratic views of perception. I hope instead to show that whatever problems lie here aren't just problems for me. We don't just have "Problems for Dogmatism," but, more generally—if there are problems—then "Problems for Credulism." That's a term I'll introduce to capture a broad range of epistemologies. It includes my

*Thanks to several forums in New York; and to audiences at Brown, Fribourg, York, Geneva, Bologna, Svolvaer, Toronto, and the Institute for Advanced Studies at Hebrew University; and to the Institute for financial support. In different ways, Maria Lasonen Aarnio, David Barnett, David Christensen, Annalisa Coliva, Philip Ebert, Adam Elga, Hartry Field, Matt Kotzen, Eugenio Orlandelli, Christian Piller, Sherri Roush, Stephen Schiffer, Nico Silins, Levi Spectre, Scott Sturgeon, Chris Tucker, Ralph Wedgwood, Jonathan Weisberg, Roger White, Tim Willenken, and Crispin Wright, made especially helpful contributions to the end result. Thanks also to Marinus Ferreira for making the diagrams.

¹There is not a single criticism, but a cluster, only some of which I will directly engage with. The provenance of these criticisms is messy. I was aware of many of the difficulties myself, and began exploring formal alternatives to Bayesianism in 2002. Cohen and Wright were at that time also pressing these complaints against me in correspondence. Shortly after, various of the criticisms were spelled out in print in: Cohen 2005; Hawthorne 2004b, at pp. 73-76; Schiffer 2004; White 2006; and Williamson 2005. Since then they've been more widely discussed. See <http://www.jimpryor.net/research/bayesian.html> for links to the work I'm aware of that deals with these issues.

own view of perception, but also *many* other contemporary epistemologies. All of them want to say things about undermining that are awkward to represent in Bayesian terms.

We'll begin by trying to get sharper about what "undermining" consists in. That will turn out to be not so easy. We will be able, though, to get sharper about what "dogmatism" and "credulism" are. Then the bulk of our discussion will critically exposit some of the supposed Bayesian problems for dogmatism. We'll see that the "problems" require some substantial assumptions about the philosophical import of different elements of the Bayesian formalism. These assumptions are widely held, but they go beyond anything that's part of Bayesianism proper, or that familiar arguments for Bayesianism establish.

2. Undermining

The way I'll use the term "defeat," it's only things you learn or get justification to believe that count as defeaters. There's an alternate use of the term, where mere facts in the world can "defeat" some epistemic status you'd otherwise have, even if you're ignorant of those facts. That's also a legitimate and interesting notion, but let's not call it "defeating."² I want to reserve "defeating" for a kind of justification you can acquire.

What are our different paradigms of defeating evidence?

Let E be the fact that Ernie tells me that his aunt's pet Precious is a bird. This supports the conclusion H, that Precious has the ability to fly. However, Orna gives me **opposing** evidence. She says that Precious is a dog rather than a bird.³

Defeating evidence need not oppose Ernie's testimony in this direct way. There are other ways to weaken the support I have for believing H, where the new evidence doesn't itself intuitively speak one way or the other about Precious's flight ability. An example: Ursula tells me that Ernie has no idea what Precious's species is; he's just guessing. She doesn't herself weigh in about Precious's real species or flight ability. I call defeating evidence of this sort **undermining evidence**.⁴

Yet other kinds of evidence are also possible: perhaps I get **refining** evidence: I learn that Precious is indeed a bird, but a predominantly flightless one, such as a penguin. Or I might learn that Precious is no bird but is capable of flight all the same. (A flying squirrel?)

²We might call it "disabling."

³Sometimes this is called "overriding" or "rebutting" evidence. Since opposing is a matter of degree, the terminology I'm suggesting is better. Depending on how the example is filled out, you may end up trusting Ernie on balance more than you do Orna; but so long as her testimony has even some impact on your credence in H, it will have opposed Ernie's testimony.

⁴Another term sometimes used is "undercutting."

We will focus on undermining defeaters.⁵

Most of us discern an intuitive kind there, but if it is a real epistemic kind, it's difficult to say in a rigorous way what makes it distinctive.

The natural first thought is that evidence that opposes the support E gives me for H will do so by *speaking for the negation* of H. Orna's evidence is an example: her testimony that Precious is a dog supports his actually being a dog, and so incapable of flight. Whereas undermining evidence, the thought goes, speaks *not* for the negation of H, but rather for some claim such as that E is unreliable with respect to H, that E should not be trusted on this matter, and so on.

This sounds like a natural way to distinguish undermining evidence from (at least some) other species of defeating evidence.

Yet it takes only a moment to notice that each of the claims we just made can be applied with some merit to the other kind of defeating evidence, too. The opposers may speak in the first place for not-H, but then by doing so, don't they also constitute some evidence against the reliability, in this context, of any evidence like E that speaks for H? The underminers may speak in the first place for the unreliability of E, at least concerning H, but then by doing so, don't they contribute to its being less reasonable for me to believe H—and so more reasonable to believe not-H—than it was before?

Perhaps the thought is significant that the opposing evidence speaks “directly” or “in the first place” for one kind of upshot, and then only by way of that for other evidence's unreliability. Similarly the thought that the undermining evidence speaks “directly” or “in the first place” for some other evidence's being in this context unreliable, and then only by way of that against the hypotheses that evidence formerly supported. But to turn those thoughts into a fleshed-out, rigorous articulation of what undermining evidence is and why it's distinctive is no easy matter. I welcome attempts to do so, but I won't pursue it further here.

A different thought is that underminers work by “screening off” the epistemic contribution E makes to H. What does this mean?

Let's at this point introduce a notion I'll call your “epistemic probability function.” This is a probability function that is intended to represent what credences you have *ex ante* or “propositional” justification to have.⁶ Of course, it may

⁵Some epistemic effects have an undermining feel, but don't obviously involve the acquisition of new evidence, so they may not be “defeaters” as I understand this term. Hartry Field reminded me of the idea that the mere articulation of a new scientific theory could “undermine” an old one, even in advance of our acquiring evidence for the new theory. I won't try to settle here how such phenomena should be categorized. My disposition, though, is to count new insights and recognitions as a species of evidence, too; so these cases *could* after all involve our acquiring defeating evidence. It'd just be evidence of a reflective, intellectual kind, rather than the sort scientists usually talk about.

⁶Some authors call this your “rational” as opposed to your actual credence function. I resist that usage because this probability function is intended to track just the character of your evidence, and the notions that informal epistemologists understand by “rational” in my

turn out that such facts aren't representable by a probability function—this is a prospect we'll return to.

Let E and H be the propositions about Ernie and Precious from before, and let U be the putative undermining proposition. Let $\text{Old}(\cdot)$ be your epistemic probability function before acquiring evidence E , and $\text{New}(\cdot)$ be your probability function after acquiring evidence E . I assume that it makes no difference whether we consider E coming first, and then U , or U coming first and then E . (We'll talk more about this assumption later.) I assume also that we can usefully talk about your probability function conditioned on the hypothesis that U , though generally you'll never get justification to be certain that U , but only more confident that it's so.

With that background in place, let's say that U screens off a contribution E makes to H when:

- (i) E supports H , that is: $\text{Old}(H) < \text{New}(H)$.
- (ii) But against the assumption of U , E does not support H : $\text{New}(H|U) \leq \text{Old}(H|U)$.⁷

Now this is not enough to give us a useful notion of undermining, for it may be that (ii) holds because U already includes all the information E would provide. Alternatively what if (ii) holds because U opposes the evidence provided by E , and does so more effectively against an evidential background that includes E : for instance, Orna might tell me not that Precious is a dog, but that if Ernie says Precious is a bird then Precious is a dog.

Let's refine the conditions as follows. We keep (i) and add:

- (iii) U doesn't speak for or against H on its own.

Let's also modify (ii) to include not just the case where U *wholly defeats* the effect of E , but also cases where U merely reduces it:

- (ii*) Against the assumption of U , E supports H less: $\text{New}(H|U) < \text{New}(H)$.

view depend on more than just that. Also, authors sometimes understand "rational credences functions" to track credences that are both held and that one has justification for holding. But I mean only to discuss the latter.

For expository convenience, I will allow myself to use expressions like "more justification to believe" to mean "justification to be more confident." But it is a substantive question, that I do not intend to take a stand on, what epistemic probabilities in the sense we're considering really have to do with acquiring more justification to have the attitude of categorical belief. Nor do I make any effort here to address what they have to do with epistemic probability locutions in natural language.

⁷Standardly the "screening off" language is reserved for cases where $\text{New}(H|U) = \text{Old}(H|U)$. Also, sometimes it is further required that E doesn't return the favor to U .

The primary home of this notion has been in discussions of causation; see for example Reichenbach 1956 and Salmon 1984.

This will count evidence that Ernie has a $\geq 70\%$ chance to be making things up as undermining, as well as evidence that he is definitely doing so.

Putting these all together, and relying on a natural interpretation of (iii), we get:

$$\text{Old}(H|U) = \text{Old}(H) \leq \text{New}(H|U) < \text{New}(H)$$

These refinements help, but they don't fully address the concerns voiced a moment ago. Like positive evidence, undermining also comes in degrees. And it can *also itself be defeated*, and *can come mixed with other opposing or supporting information*. It may be no easy matter to disentangle these different evidential components, when they do come mixed together.

Here's a series of examples that illustrate that possibility, and obstacles it poses to analyzing undermining in terms of screening-off.

In each of the examples, I'll possess some auditory evidence, which will be undermined by olfactory evidence of smoke. I'll assume background knowledge that smoke tends to make me have random auditory hallucinations more frequently. (Perhaps it's opium smoke?)

In the first case, the evidence that the smoky smell threatens to undermine is my auditory experience of a radio. On its own, that experience would give me some justification to believe a radio is nearby. In the presence of the smoky smell, though, this experience cannot reasonably be relied on to the same degree. I now have more justification than before for thinking the auditory experience is hallucinatory. It doesn't matter *how much* less justification to believe there's a radio nearby this results in; it's enough for our purposes that it results in some. Of course, I can't be *certain* that there is smoke, much less that I've auditorily hallucinated. What I'm hallucinating might instead be the smoky smell. But I needn't be certain, for the evidence of smoke to have some undermining effect on my auditory justification that there is a radio nearby.

In the second case, the evidence that the smoky smell threatens to undermine is instead my auditory experience of singing birds. We will assume background knowledge that when there really is smoke present, birds flee. They don't stick around and sing. Now, in this case too the smoky smell contributes to making it reasonable for me to rely less on my auditory experience. However, what's different in this case is that the smoky smell *also* on its own speaks against the same proposition my auditory experience seemed to support. That is, the smoky smell on its own testifies against there being any birds present. So what we have here is that the smoky smell *both undermines and opposes* my auditory evidence for birds.⁸

In the third case, the evidence that the smoky smell threatens to undermine is instead my auditory experience of an approaching fire engine. On its own, that

⁸For another example like this, see the case of the butler and the maid in Weisberg ms.

auditory experience seems to support the hypothesis that firemen are nearby. However, in this case too, the smoky smell calls for me to be more guarded towards my auditory experiences. Perhaps this too is just another smoke-generated auditory hallucination. On the other hand, my olfactory experiences themselves seem to improve the likelihood that firemen are nearby. Where there's smoke, there's usually firemen. There may be a complex interaction between the undermining effect my smoky experiences have on my auditory evidence, and the independent olfactory evidence for firemen they put in its place—for there actually being firemen there to be heard presumably makes it again somewhat less likely that I'm hallucinating. We haven't said enough to know what equilibrium this leads to. What's important is that here the smoky smell to some degree undermines one body of evidence for conclusion H, while at the same time itself providing *other evidence for H*.⁹

I call these last two cases **mixtures**: in the one case, a mixture of undermining and opposing, and in the other case, a mixture of undermining and supporting. They are the shaded regions in figure 1.

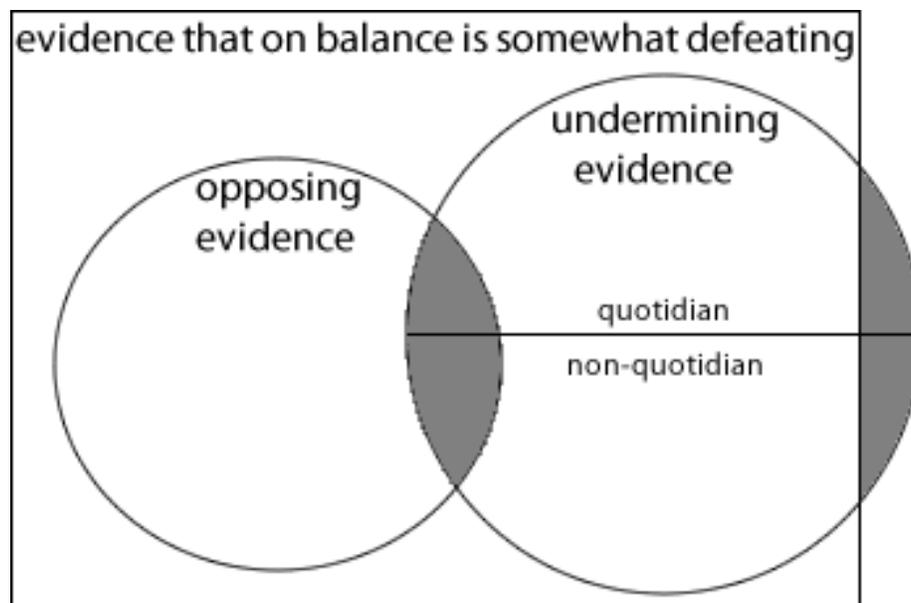


Figure 1:

I'll explain the contrast between quotidian and non-quotidian undermining later.

Some of the time, as perhaps in these examples, we may be able to separate out different considerations in a mixed body of evidence, which point in the different directions. However, I don't think we're in any position to assume *it will always*

⁹Silins forthcoming gives some examples of this sort.

be so. Mixed bodies of evidence may not always be easily decomposed; they may for all I know sometimes be in principle un-decomposable.

Mixed evidence of these sorts makes difficulties for the project of explicating undermining in terms of screening-off. A screening-off test seems like it would, at best, capture what's going on in cases of **pure undermining**, with no mixture of additional opposing or supporting effects. It won't capture cases where the combination of U and E on balance *increases* the likelihood of H (because U adds more support of its own for H than it diminishes E's). Because of condition (iii), which aimed to exclude *mere* opposing evidence, our screening-off test can fail to capture cases where U *both* undermines and opposes E's support of H. We should hope for an account of the distinctively undermining evidential effect as it shows up in all of these cases.¹⁰

Some have tried to do better at spelling out undermining in probabilistic terms.¹¹ We won't pursue those attempts here, either. It's enough that we've seen some of the initial difficulties. In the rest of this essay we'll see other sorts of awkwardness for reconciling Bayesianism with some intuitive views about undermining. For these purposes, I propose we just proceed with our intuitive, as-yet-unanalyzed understanding of undermining.¹²

¹⁰David Barnett offered the following example, where your belief is intuitively undermined but *no* change of credence in the relevant proposition is called for. You're at a crossroads, inquiring the way to Camelot. Your informant tells you it lies to the right. You know your informant is either a knight, who is perfectly reliable, or a knave, who is perfectly anti-reliable, or a fool, who answers at random. You initially estimate he's 45% likely to be knight, 45% knave, 10% fool, and so you suspend judgment whether Camelot really does lie to the right. Still, you do have some evidence to think so; it's just balanced by equally weighty evidence to think not. Next you learn that your informant is the fool. Intuitively, some kind of undermining has now taken place. Earlier the informant's testimony seemed 50% likely to be accurate, and 45% likely to be accurate because reliable. Now your evidence still seems 50% likely to be accurate, but not at all likely to be reliable. Because of the way the case is set up, no change in your credence about Camelot's direction is called for; but your grounds for believing Camelot lies to the right seem intuitively weakened from what they were before.

This case invites the idea that although undermining your justification for H may not be a matter of how your credence *in H* changes, it may be a matter of how your credence in some other proposition changes—perhaps a proposition about the reliability of your evidence for H. But read on, and see the variety of undermining considerations we display in section 4, below. Then try and say specifically what other proposition it is we can analyze the undermining of your justification for H in terms of. Also, would you insist that no subject's justification for H can be undermined who lacks defined credences for propositions about reliability, their own epistemic status, and so on? I acknowledge that the notion of "defined credence" we're working with is an epistemic one, not a psychological one. But I would still be reluctant to think every subject who is vulnerable to undermining must have defined credences for these things.

¹¹See for example Kotzen ms.

¹²I'll mention briefly a last strategy for explicating undermining evidence. John Pollock was the first I know to discuss the phenomenon of undermining in a sustained way. (For a survey of his views, see Sturgeon forthcoming.) Pollock called this phenomenon "undercutting", and proposed that U undercut E's support of H just in case U was evidence against the conditional "E wouldn't be true unless H were." Interpreting this is not easy, because it involves an "unless" construction, the subjunctive mode, and the notion of evidence against a conditional—each of which is tricky in its own right. (Also, it's not clear how seriously we should interpret its similarity to Dretske 1971's definition of a "conclusive reason.") I think

3. Dogmatism and credulism

When I began several years ago to use the term **dogmatism**, I meant it to be the view that immediate justification exists even where it might be undermined by skeptical (or mundane) defeaters that one has no epistemically antecedent grounds for ruling out. That is, justification is sometimes both immediate and underminable.

More specifically, what I mean by **immediate justification** is some amount (it matters not how much) of prima facie justification *to* believe something (that is, prospective or *ex ante* justification) that does not even partly come from—or, using other language, is not even partly constituted by—your having justification to believe something else.

Justification is on the other hand **mediate** when it does in part come from, or is in part constituted by, your having justification to believe something else. So understood, your justification for some belief might be mediate even if the belief was arrived at spontaneously, not via any chain of explicit reasoning. And even in cases where your belief *was* explicitly (and competently) inferred, it might not be *based on* all the upstream considerations whose justification entitled you to have it.¹³

I've advocated a view about perception that emphasizes the possibility of immediate but underminable justification, and many have come to associate the term “dogmatism” with that particular view. I've also said things favorable about “Moorean” arguments against skeptical hypotheses (though also some things unfavorable). And again, some have come to associate the term “dogmatism” with sympathy for such Moorean arguments. However, it is a substantive claim

the proposal is most plausible when the relevant conditionals are understood epistmically. If we understand them truth-functionally, then the proposal seems to be that something undercuts E's support of H iff it's evidence against $E \supset H$. Both directions of this are odd: why couldn't there be underminers which make E and not-H *less* likely (though more likely than they make E and H)? And if Orna tells me that Precious can't fly (not-H) but that Ernie will tell me it can (E), must that be counted an underminer of E? Understanding these conditionals counterfactually brings other difficulties. So as I said, the proposal seems most plausible if the conditionals are understood to express something epistemic, along the lines of “If E, then likely also H,” or “If E, then it might be that not-H.” (Pollock himself glosses his conditionals as “E does/doesn't guarantee H”; but these are open to more-or-less the same range of interpretation.)

And yet, what does it mean for U to be *justification for* (or against) an epistemic conditional? On some views, such conditionals don't have truth-conditions and so perhaps aren't legitimate targets of justification. Even if that thought is mistaken, and it really does make sense for there to be justification for or against epistemic conditionals, I submit it's not a sense we yet have any good understanding of. So even if Pollock's criterion ultimately does prove to be correct, I don't think it's a suitable place from which to begin our inquiries.

¹³It's a substantial question whether the grounds of a competently-based belief need to include all the considerations justification for which make up the justification you have for that belief. If they do, and we subscribe to commonly-held views of which considerations those are, and commonly-held understandings of the basing relation, it's doubtful that very many beliefs are competently-based.

that these go together, and should not be assumed as a matter of definition.¹⁴

It's not up to me how others will use the term "dogmatism." But I will use it to name the general thesis that justification is sometimes both immediate and underminable. It doesn't include by definition any commitments about Moorean arguments. Neither does it include by definition any commitments about *why* some body of justification is immediate. Some of us favor an internalist view of perceptual justification, that ascribes the justificatory power to the quality of our perceptual phenomenology.¹⁵ But that's just one species of dogmatism. Other stories are also possible. And those stories need not always be internalist. A reliabilist or a disjunctivist about perceptual justification can also claim that justification to be immediate yet underminable.

Finally, "dogmatism" in the general sense used here isn't a thesis specifically about perception. Perhaps you deny that perceptual justification is immediate; but you think your knowledge of what you intend to do is often immediate—yet also underminable. Or you think mathematical justification is sometimes immediate—yet also underminable. My own inclination is to expect all justification to be underminable, and I'll try to persuade you to think so too. So if you think we ever have immediate justification about anything, you should be a dogmatist, too.

Many philosophers are already dogmatists in this sense, though they may never have applied that label to themselves. However, much of what we're going to discuss bears also on an even more inclusive group, which has not before been named. I will call this group **the credulists**.¹⁶ We will look at what defines the group more carefully in the next section. For the moment, here is a quick gloss. Credulists think you can be justified in believing H, in a way that would be undermined by evidence for U, without antecedent justification to believe not-U needing to be a constitutive part of your original justification for H. We can sum this up in the slogan: **Your justification for believing H is vulnerable to being undermined in ways you didn't need to antecedently rule out.** A dogmatist adds the additional commitment that this is because *no* antecedent justification to believe anything was part of your original justification to believe H. A credulist need not be a dogmatist, though. They are allowed to think that antecedent justification to believe some things was part of your justification to believe H. Maybe even antecedent justification to rule out some underminers. But there are *some* propositions like U that would undermine and *didn't* need to be antecedently ruled out.¹⁷

¹⁴Silins 2008 argues for the first without the second. See also Wedgwood forthcoming; Neta 2010; and Kotzen 2012.

¹⁵Besides myself, see also Huemer 2001a, 2006, 2007; Tucker 2010; and Chudnoff 2011.

¹⁶This is just an arbitrary label; there's not supposed to be any tight connection to the folk meaning of "credulity." I wanted a term that sounded reminiscent of "dogmatism" but more moderate and so more inclusive.

¹⁷A comment on locutions like "In order to be justified in believing H, you need/require antecedent justification to believe not-U." When I say this, I always have in mind the claim that the antecedent justification is included in your justification to believe H. There's another

In the next section, we will call underminers of this sort **non-quotidian**. So credulism is commitment to the possibility of non-quotidian undermining.

In later sections, we will see difficulties for representing non-quotidian undermining in either Classical or Jeffrey Bayesian terms. These difficulties won't come from the formalisms on their own, but from the combination of the formalisms with a popular set of interpretive assumptions, which I will identify.

It shouldn't be surprising that epistemic effects belief in which make one a credulist should be difficult to represent in Bayesian terms. For consider the following:

Hypotheses H1 and H2 are logically equivalent, and as a matter of fact you do justifiably give them equal credence. However, your own mental states are not transparent to you. Suppose you now acquire evidence that your credence in H2 is in fact lower than your credence in H1. It is debatable what effect such evidence should have, but there is a presumptive case that it would put *some* kind of pressure on you to raise your credence in H2 and/or lower your credence in H1. Refusing to alter your current credences would seem unreasonably unresponsive. True, the evidence you've just acquired is misleading, but it's part of the case we're imagining that you're in no position to know this.

Or:

Hypotheses H1 and H2 are logically equivalent, and as a matter of fact you do justifiably give them equal credence. However, their logical relationship is not transparent to you. Suppose you now acquire evidence that H1 is in fact logically stronger than H2, and that the possibility of H2-but-not-H1 should have a positive credence. Here too, there is a presumptive case that the evidence would put *some* kind of pressure on you to raise your credence in H2 and/or lower your credence in H1. Here too, the evidence is misleading, but you're in no position to know this.

(perhaps more literal) reading of the locution, where it only says that having some antecedent justification to believe not-U is a necessary condition to have justification for H. I do acknowledge the difference between these claims (see note 33 of Pryor 2000); and Silins 2008 argues for the importance of separating them—partly on the basis of the Bayesian issues we'll be examining below. (See also his forthcoming; and McGrath forthcoming, section 3.) However, I will neglect the difference in this discussion, because the dialectic is already very complex. “Neglect” not “ignore”: I won't assume there is no difference. But it may be that some of what I say needs to be revised when we attend to the difference more carefully.

Neta 2010 argues for this distinction and for others, some of which I also acknowledge but describe differently. My dogmatism about perception is what he calls “mediate liberalism.” I'm not sure I recognize any difference between his “simple liberalism” and his “inferential liberalism.”

It's already commonly granted that Bayesian formalisms aren't straightforwardly able to represent epistemic effects of these sorts. But many cases of non-quotidian undermining seem, at least to this author, to be cut from the same cloths.

If it's in fact true that non-quotidian undermining can't be represented in Bayesian terms, what should be our response? Should we conclude that non-quotidian undermining is impossible? Should we rein in our explanatory ambitions, and say the formalisms only apply where we can idealize in such a way that the possibility of non-quotidian undermining can be ignored? Or should we give up some of the interpretive assumptions that make non-quotidian undermining unrepresentable?

We will return to these choices.

4. Credulism more carefully

To get a better handle on what credulists believe, let's consider the following example.

You have the evidence E, that a certain barometer is falling. E together with other things you are justified in believing—for example, about reliable connections between that barometer and the upcoming weather—justify you in believing H, that there will soon be rain.

Though our focus is prospective or *ex ante* justification, for some of the variations I want to consider next, we should suppose that this transition from E to H is an inference you've explicitly drawn.

So, how can the justification you have for H be defeated, while leaving your original evidence in place—and intuitively without directly speaking for or against H?

The most quotidian way for this to happen is for the defeating evidence to oppose some auxiliary hypothesis, where it's also plausible that you needed antecedent justification for that hypothesis, to have the justification you did for H. In our example, evidence that your barometer is not reliable about the weather would be of this sort.¹⁸

There are things one can fuss about here. Perhaps we shouldn't say it was ever E itself that supported H, but only the combination of E and the hypotheses that connected it to H. If so, does this really deserve to be called a case of undermining? Undermining requires that what was evidence for H is still available,

¹⁸As mentioned before, your needing antecedent justification for an auxiliary hypothesis needn't mean it had to be among the *grounds on which you based* your inference that H, for that inference to be competent.

but supports H less. This instead is a case where you start with evidence E-and-connecting-hypotheses to believe H, and new evidence takes that old evidence away. You're then no longer (or less) justified in believing E-and-connecting-hypotheses.

It doesn't matter to me whether these quotidian examples do deserve to be called examples of undermining. What interests me is that many of us think there are possibilities of undermining *other than* what's just been described. That is, **non-quotidian** forms of undermining.

For example, a dogmatist can think your justification *for E itself* is underminable, but immediate, so that there are *no* hypotheses you needed antecedent justification for, to have the justification you did for E.

Credulists more generally take the same attitude towards some underminer that the dogmatist does: it's a way in which your justification is vulnerable, that you didn't need antecedent justification to believe wasn't realized. But a credulist needn't insist that *no* antecedent justification was involved. In our example, your evidence for H did require antecedent justification to believe the barometer was reliable. But arguably it *didn't* require antecedent justification to believe the higher-order claim that E plus evidence of reliability *are enough to justify* belief in H. Yet the credulist may think that evidence that these *aren't* enough would justify you in being less confident, or in some way more guarded, towards H.¹⁹ Such evidence may take the form of a compelling philosophical argument that there is no justified belief about the future. Or that, though there may be such, it is not available to subjects with the cognitive shortcomings you know yourself to have. Shouldn't evidence for those claims undermine your confidence in H at least somewhat?²⁰

Suppose we concede your original justification for H *does* have to include antecedent justification to reject the conclusions of those arguments.²¹ Then consider instead evidence that what was *in fact* your antecedent justification for doing so, wasn't. Shouldn't evidence for *that* undermine your confidence in H at least to some degree? And did you need antecedent justification to reject *it*?

In all of these cases, I am supposing that your original evidence really did support H; that enables us to sidestep questions about how to deal with illusions

¹⁹I count such cases as examples of undermining, though there may be interesting differences between them and cases of undermining that aren't higher-order.

Some philosophers will urge that degrees of confidence don't by themselves provide the right structure to model the relevant guardedness. I have some sympathy for that, but won't pursue it here.

²⁰Cases like this, and the ones that follow, have been much-discussed in the disagreement literature, especially by David Christensen. See Christensen 2007, 2008, 2010a, 2010b, 2011; Elga ms; Feldman 2005, 2006, 2009; and Kelly 2010, esp. section 4. See also Huemer 2011 and Schechter forthcoming.

²¹This is in the vicinity of what I call "Inferential Internalism" in Pryor 2001. I take that term from one of its proponents: see Fumerton 1995, Chapter 3. Alston opposes the view in several essays in Alston 1989. For more recent work on this issue, see Tucker 2012 and the work he references. Boghossian 2003 argues that for deductive inference, simple forms of Inferential Internalism and Inferential Externalism are both wrong.

of support. Hence, the undermining evidence is evidence for a false claim. But that surely is no obstacle to your acquiring it.

Some philosophers I've discussed these issues with have resisted the idea that philosophical arguments for false conclusions can provide justification for those conclusions. They seem to me to have unrealistic conceptions of philosophical justification. But we needn't fight about that. The undermining evidence needn't be provided by any mistaken philosophical argument, itself. Rather, let it come in the form of empirical evidence about *how your philosophical mentors assess* arguments you don't yourself understand, or haven't yourself seen.

Other philosophers I've discussed these issues with have complained about this defeating evidence being directed in the first place at questions like whether you're justified in believing H, rather than H itself. Of course I agree those are separate questions. It may even be that we can not expect a straightforward correlation between your epistemic position with respect to the one and your epistemic position with respect to the other.²² But that is compatible with evidence for the higher-order claim having *some* defeating effect on your justification for the lower.²³

Another kind of underminer would be (misleading) evidence that no claims about the future are *true*. Shouldn't that put some pressure on you to *withhold belief* about the future? Must your justification to believe H have included antecedent justification to disbelieve any such metaphysical evidence?

One can have credulist views, not only about vulnerabilities in your justification to believe some *hypotheses*, but also about vulnerabilities in *transitions or inferences* you make. You might be entitled to less confidence in H because of evidential challenges to your *move* to H, from premises that do in fact support it. The challenges we considered testify that such moves are inappropriate even in the ideal. There can also be challenges to your actual performance of the move: that is, (misleading) evidence that you have not reasoned competently in inferring H from E.²⁴

It is debatable what effect such evidence should have, but there is a presumptive case that it should put *some* kind of pressure on you to be less confident in H on that basis. This is compatible with there *also* being some pressure on you *to* believe H on that basis, since the basis does in fact support H. But so long as the first pressure also exists, and is not wholly trumped by the second, we would have a case where your confidence in H was undermined to some degree.

Now, in order for you to have the justification you originally did for H, did you need, as a constituent, antecedent justification to believe you were inferring H competently?²⁵ Many epistemologists will be reluctant to think so. That would

²²Williamson 2011 argues that knowing P is compatible with the epistemic probability that you know P being arbitrarily low. See also Williamson ms.

²³See here also the end of Kotzen 2012.

²⁴Willenken 2011 focuses on defeaters of this sort.

²⁵This seems to me different from, and much less plausible than, Inferential Internalism.

make considerations about your epistemic biography essential parts of bodies of evidence where we would not have expected to find them. There'd then be no wholly mathematical or a priori justification, since evidence about your competent performance would have to be part of the story too, before anything got to be justified.²⁶ Nor could there be any wholly historical justification.

But for the sake of argument, suppose we concede that. Let it be agreed that any time you believe H on the basis of E, part of what justifies the belief in H must be antecedent justification to believe you're performing that inference competently. Then consider instead what should be the effect of (misleading) evidence that *you are not properly taking account of* the evidence of your competent performance, in believing H on the basis you do. Again, it is debatable what effect such evidence should have, but there is a presumptive case that it should put some kind of pressure on you to be less confident in H on that basis. Now was *that* something you needed antecedent justification to disbelieve, as part of your original justification for H? If so, we can keep going.

The credulist thinks that at some point, we will find some way in which your original justification is vulnerable to being undermined, without your original justification needing to include antecedent justification that that vulnerability isn't realized.

I can see only two non-credulist alternatives. One says, No, you really did need a tower of antecedent justification as part of your original support. And every underminable part of that tower needs to be supported by its own antecedently justified tower, and so on ad infinitum.²⁷

The other alternative is a specific kind of "externalism." At some point—perhaps at the very first step—it would say, if you are in fact doing everything properly then evidence that you're not should have no defeating effect. Moreover, it wouldn't say this just once, but *everywhere* the prospect of a not-antecedently-ruled-out vulnerability arose. Don't worry, this view would say. You don't need an infinite tower of antecedent justification to be justified in believing H. But that's not because the credulist is right, and some vulnerabilities don't need to be antecedently ruled out. Rather, it's because these are not really vulnerabilities of the operative sort. If in fact you reasoned properly in such-and-such a respect, your confidence shouldn't be threatened by evidence that you didn't.²⁸

Both of these alternatives are intelligible. But they seem to be minority views among epistemologists I interact with. I expect that most epistemologists will

²⁶Silins forthcoming also mentions this threat.

²⁷This couldn't even be given a coherentist spin, because the view we're envisaging says that every underminable part of a tower is supported by *antecedent* justification to believe other things. This is a notion of epistemic not temporal priority, but presumably any notion of priority will exclude cycles.

²⁸This is reminiscent of the view of disagreement defended in Kelly 2005; though Kelly himself explicitly *refrained* from claiming that higher-order evidence has no defeating effect (see his section 6). And in Kelly 2010, section 4, he argues against that claim.

be reluctant to hold them, and will instead turn out to be credulists.²⁹

To summarize: Dogmatists think some vulnerabilities don't need to be antecedently ruled out, because the justification that's vulnerable is immediate, and doesn't include *any* antecedent justification. That's one way to be a credulist. Other ways to be credulist say that, even where the justification in question requires *some* other antecedent justification, it doesn't require antecedent justification against *all* of the ways—perhaps arbitrarily higher-order ways—in which it may be undermined.

I have no master argument that it's impossible to be a credulist and also a Bayesian. However, as we'll see, there are widespread assumptions about how to interpret aspects of the Bayesian formalism—what their philosophical “cash-value” is—that seem to get in the way of doing so.

5. Some assumptions and ambitions

The first assumption to consider is:

Assumption-1 Getting more justification for H coincides with raising H's probability.³⁰

This has been challenged. Geoff Pynn imagines you start out with higher credence in H than you should have. Say you should have a credence of 0.5, but in fact you have a credence of 0.9. Then you acquire evidence in the light of which you should have a credence of 0.6. At this point, should your credence go up? Well, not up from 0.9. It should go down, because it was too high to start with.³¹

There are interesting issues here about the relation between the attitudes it'd be normatively correct for you to have, and the attitudes you actually do have. As I discuss elsewhere, I do not think those facts are insulated from each other.³² However, in the present context, we are understanding talk of your “credences”

²⁹The view that Willenken 2011 calls “liberalism” is roughly equivalent to credulism (modulo the issue mentioned in note 17).

The “defeasibility argument” that Silins considers in his forthcoming is, in one of its initial variations, an argument that all underminers are quotidian (and so against credulism). Silins criticizes both that argument and some more restricted variations.

³⁰White 2006, note 10 points out that his argument against dogmatism needs only a more cautious version of this, which says that getting more justification for H is incompatible with *lowering* H's probability. That is true, but I have never heard any challenges to the bolder assumption that would not also bear against the more cautious one, so I'll just discuss the bolder, simpler assumption.

³¹Pynn forthcoming, in section 3. See also Enoch 2010, section 5; and the example Christensen 2011 calls “Wrong and Wronger.” (The name is Christensen's, but the case comes from Kelly 2010.)

³²See Pryor 2004; Pryor forthcoming; and Pryor ms.

and your “probability function” to be confined to the attitudes it is reasonable for you to have, prospectively or *ex ante*, regardless of what attitudes you do in fact have. So in the given example, your epistemic probability for H—the credence your epistemic position calls for—does go up, from 0.5 to 0.6, regardless of your actual doxastic mistake.

A different concern with Assumption-1 is that our informal notion of justification is in the first place rooted in thoughts about prima facie support, whereas the Bayesian works with the resultant sum of many different justificatory pressures. For example, you may initially estimate the probability that the notorious pet Precious is a bird at 0.6, and that it’s a penguin at a minuscule 0.01. Peter then testifies to you that Precious is in fact a penguin. You may trust Peter enough that the probability Precious is a penguin climbs substantially, say to 0.4. But the prospect of a pet penguin strikes you as so odd that you also suspect Peter may have seen some other, non-bird animal that he confused with a penguin. So you’re somewhat less sure now that Precious is a bird—perhaps only 0.5. You acquired justification to believe Precious is a penguin. Did you fail to acquire justification to believe it’s a bird? I’m reluctant to say so. I’d rather say, Yes you did acquire *some* justification to believe that, it’s just that all-things-considered, you’re left in a position where it’d be, on balance, reasonable for you to be less confident that Precious is a bird than it was before. Our simple informal epistemic notions more closely track the prima facie notions. It takes more words to describe the all-things-considered facts.

I explained dogmatist views as in the first place concerned with what’s required to have prima facie justification. Still, it’d be a real disappointment to hear that, despite our experiences giving us prima facie immediate justification to believe our environments are as they seem, we’re never all-things-considered justified in being more confident of that than we are that we’re in some skeptical scenario. So dogmatists will want to say that their prima facie justification *can* in some cases amount to ultima facie justification, too. They won’t just turn their back on the Bayesian and say, we’re talking about different epistemic notions. But it is worth keeping in mind that the theories aren’t *directly* working with the same epistemic quantities. This point will come up again later.

I will not directly contest Assumption-1 any further, though some of what we’ll say below should inform our later assessment of it.³³

We’ll work our way to the next assumption by starting with a concern a fan of immediate justification may have about how to represent her view in the Bayesian framework. The Bayesian says that epistemic changes are (when reasonable) always a matter of updating on propositions. But why should we think that acquiring immediate justification in the way our informal theorist believes possible corresponds to *that*? For example, a core claim of my dogmatist view of perception is that it’s *merely having* certain experiences that justifies you in believing you have hands—not *learning that* or getting reason to believe that

³³For further discussion, see Achinstein 2001, Chapter 4; and Kung 2010.

you do have those experiences. And we can expect dogmatist views generally to have this character. They'll propose some kind of epistemic situation S, and they'll say it's merely *being in S*, not having justification to believe you are, that justifies you in believing other things. Wouldn't modeling such epistemic changes in terms of updating on a proposition distort what the dogmatist thinks is going on?

This is not a simple issue. I'll sketch a first response now, but shortly we'll see this worry re-arise in a different form.

The first response says: That's OK, this much of what the dogmatist thinks is going on can be reconciled with the Bayesian model. For on these dogmatist views, there will be some first propositions that getting into S makes it reasonable for you to believe. So we can say it's *those* propositions that you reasonably update on. That is compatible with our having an extra-formal, dogmatist story about *why* you're in a position to update on it.

All right. But let's think more carefully about *which* propositions it is that you update on. The response we just heard suggested:

Assumption-2 What you update on coincides with what you're immediately justified in believing.

This has two directions: first, that your probability function updates on E only when you've acquired immediate justification for E.³⁴ The other direction is that changes in your credence in H that aren't themselves updates on H, but are rather a function of your updating on something else, cannot represent acquirings of immediate justification to believe H, but only the acquirings of mediate justification. If a theorist doubts the existence of any immediate justification at all, she might uphold the second of these without the first.

Assumption-2 looks natural, but it should not be assumed uncritically. Recall that immediate justification is in the first place defined for prima facie justificatory contributions, whereas what you update on will be the ultima facie, net sum of different such pressures. So what you update on might not be *any* proposition you then acquired prima facie justification for, much less immediate such justification.

Moreover, alternatives to Assumption-2 are intelligible.

On some views, our probability functions just simply *evolve*, and facts about what propositions we've "updated on" are *reconstructions* from that evolution. We are not given as independent facts that now you *should* update on this proposition, and now on that. On such a view, I see no reason to expect we could read off facts about immediate justification from what propositions you've

³⁴Remember, these "updates" concern how the credences you're *justified* in having change. Of course unreasonable subjects may as a matter of descriptive psychology change their doxastic attitudes in ways that don't correspond to any justification they've acquired.

updated on. (On the other hand, I'm not sure such views are hospitable to the possibility of immediate justification, in the first place.)

Another view says that at a given time you become entitled to certainty that you're having handlike experiences, and reasonable confidence (but not certainty) that you have hands, and the second of these *does not epistemically depend on* your justification to believe the first. You become immediately justified in believing each. (Since we're credulists, we may allow that evidence *against* the first would tend to undermine your justification for the second.)

This view seems coherent. Assumption-2, however, implies that it's only the claim about your experiences that you could have acquired immediate justification for.

If we're going to engage with dogmatists, then we should be able to *at least try* to represent justification that's both immediate and underminable. But if it's immediate, then Assumption-2 tells us you update on it. And in the Classical Bayesian framework, when you update on something it becomes maximally justified and you can never take it back. No future evidence can defeat it. When we turn to Jeffrey Bayesianism, later, the formalism will permit you to update on hypotheses without becoming certain of them. Perhaps that will prove more hospitable to the dogmatist. But is there any way to make *Classical* Bayesianism (at least initially) more accommodating of what the dogmatist thinks happens?

Perhaps we should just give Assumption-2 up. We could suppose that updating on the claim that you have handlike experiences is compatible with your having acquired immediate justification to believe bolder claims. We will see White take this approach. An advantage of such views is that it will be *possible for the subject to recognize* she's in a situation that gives her immediate justification to believe H, without thereby committing herself to H, because it has been undermined.

Alternatively, we might try appealing to "Popperian" conditional probabilities, which are defined even for conditions that have an unconditional probability of 0. So it could be that your $\text{New}(H)$ is 1, yet you still have well-defined, non-trivial values for $\text{New}(.|U)$, where U is incompatible with H. This is not the domain where these tools are usually used, but it may be a viable application of them. On the face of it, this would permit us to update on the immediately justified propositions, but allow them later to be defeated.

But how should we think about cases where you acquired the defeating evidence *first*, before your immediate justification for H? Intuitively, that shouldn't make a difference to what your final epistemic position is. We will return to this intuition later, and discuss it more carefully. For the time being, assume it's right. But if, when you get the defeating evidence *second*, the credence you're justified in having in H goes lower than 1, then in the case where you get the defeating evidence first and the immediate justification for H second, you shouldn't there be Classically updating on H. Your credence in H should in that case not go to 1. Yet H is what you acquired immediate justification to believe—it was

just preemptively defeated. So even if we want to use “Popperian” conditional probabilities to let you update on H in some cases, we’ll still end up denying that Assumption-2 always holds.

Recall that the Bayesian formalism itself is completely silent about when subjects should update on which propositions. Plantinga tells a useful story of a mountain climber whose beliefs “freeze” at a certain moment, though his sensory experiences continue to evolve like everyone else’s. His friends carry him home, but he continues to believe he’s sitting on a ledge watching a hawk glide below him. Yet he’s not hallucinating: he has exactly the perceptual phenomenology any of us would have, at home in his bed. It’s only his beliefs—including his beliefs about his experiences—that have frozen; not the experiences themselves.³⁵ Plantinga’s point is that if the subject’s beliefs were coherent just before the freeze, then so far as their internal relations go, they’ll continue so—as long as no new beliefs are added or removed. But surely this subject manifests *some kind* of epistemic defect. The defect is in the correlation between what’s happening in the world and how his beliefs are updating. This is a matter that can go epistemically better or worse. The Bayesian formalism doesn’t itself say anything about it. Instead, it just *starts* from the point where it’s given that the subject now does (or should) update on such-and-such.

By itself, that’s no weakness in the formalism. It’s just one of the things the formalism doesn’t try to explain.

And against that background, we can understand the view sketched a moment ago like this: when the subject has handlike experiences and reasonably low credence in undermining hypotheses, then it’s epistemically appropriate for her to update on bolder claims, like the claim that she has a hand (or perhaps the claim that she sees herself to have a hand). When she has the same experiences, but reasonably higher credences in undermining hypotheses, she should instead update on more cautious claims about handlike experiences. So far, nothing here conflicts with Bayesianism. It just goes farther than Bayesianism does and says something about when the world makes one update appropriate rather than another—a matter about which Bayesianism itself is silent.

Jeffrey Bayesianism promises to be an even more hospitable setting for this strategy, too, because it can model you partly updating in the bold way and partly updating in the cautious way. But despite the crudeness of the present version, it should still be intelligible.

As I said, this view gives up on Assumption-2, for even in cases where the subject has high credence in the underminers, she still acquires immediate justification to believe the bold claim. It’s just that that justification is preemptively defeated. Yet on this view, in such cases the bold claim is not what the subject updates on.

This view also gives up the idea that there is a single proposition P such that acquiring a given piece of immediate justification always coincides with updat-

³⁵Plantinga 1993, at p. 82. See also Feldman 2003, at p. 68; and Christensen 1992, note 1.

ing on P. So the worry we voiced earlier would be vindicated: acquiring that immediate justification *isn't* correctly modeled by any one update.

Thirdly, this view abandons the following idea:

Assumption-3 The negative effect of undermining evidence, even when it purports to be immediate justification that's undermined, should be represented *inside* our formal models, rather than “off-stage,” in terms of when the world supplies the model one input rather than another.

Though I call this an “Assumption,” it is really more of an explanatory ambition. And it's quite a strong ambition. Recall that the examples of introspective and logical opacity we discussed before are already conceded to be places where it fails. Those involved epistemic effects that these formal models are already acknowledged not to represent.

I think that abandoning Assumption-3 for non-quotidian underminers is in fact a theoretically fruitful strategy. But for this discussion, we will try to see how far we can get without abandoning it.³⁶

I will identify three other assumptions when they become relevant in the discussion below.

6. Threats from Classical Bayesianism

Now we're ready to assess the alleged Bayesian problems for my dogmatist view of perception. As we'll see, if there are problems here they aren't specific to perception—and they may not be specific to dogmatism either, but be issues that all credulists need to sort out. Additionally, any problems to be found here depend on one's making certain choices about the interpretive assumptions. We've already seen a range of options for what a dogmatist might say you update on, when you acquire immediate justification to believe H. White and other Bayesian complainers against dogmatism rely on a specific one of those options. We'll soon encounter yet another interpretive assumption that is essential to their complaints.

White's discussion contains several threads. One of these concerns “bootstrap-ping”; I will not attempt to sort those issues out here. A second concerns justification or suppositions about what justification you will have in the future; we will take these issues up in the next section. A third thread concerns “Moorean” arguments like the following:

E. I am having experiences as of hands.

³⁶Section 4 of Christensen 1992 emphasizes the costs of abandoning this Assumption.

Hand. So I have a hand.

If I have a hand, I am not a handless brain in a vat being fed illusory experiences as of hands.

Good. So I am not a handless brain in a vat being fed illusory experiences as of hands.

It will be convenient to have a short label for the negation of Good, which I'll dub:

Bad. I *am* a handless brain in a vat being fed illusory experiences as of hands.

White claims that by Bayesian lights, having experiences as of hands should make it *more* likely, not less, that you are a handless brain in a vat having such experiences. For the brain in a vat hypothesis we're considering, Bad, *entails* that you'd have those experiences. Hence, when you update on E, the probability of Bad should go up, and the probability of its negation Good should go down. It'd be strange to say that such reasoning increases or contributes to the justification you have to believe Good, when the justification you acquire for the premises makes the probability of Good *go down*.³⁷

White is here criticizing those of "Moorean" sympathies, who think this kind of reasoning *can* contribute to one's justification to believe Good. I am such a philosopher. However, my project today is to speak on behalf of all dogmatists and credulists, not just those who share some specific and controversial views with me. So I won't speak to this complaint directly.

Note that White is assuming here that acquiring the immediate justification my view concerns itself with—which happens when you have handlike experiences—should be modeled by your updating not on Hand but on E. Hence, White is rejecting what we called Assumption-2. He's supposing that a charitable representation of what the dogmatist thinks happens can involve you updating on something other than what you (allegedly) acquire immediate justification for. If he's not to have already begged the question against the dogmatist in modeling things this way, then the mere fact that your credence in Hand goes up as a function of your updating on something else shouldn't be understood to *already* mean that the justification you acquired for Hand was epistemically posterior to justification you acquired to believe something else, and hence non-immediate.

Attend to this point well, for it will bear on what happens later. If we thought that modeling the situation in this way *did* mean that the justification you acquired for Hand was epistemically posterior to something else—as we would if we held Assumption-2—then White's "argument" against the dogmatist would already here be concluded. It would consist in his assertion that when you have

³⁷See note 30.

handlike experiences, you don't get immediate justification to believe Hand. All the subsequent details about probability would be irrelevant.

So in charity to White, we assume he is charitable to his target, and does not take this aspect of the model to have that significance. Later, though, we'll see he seems to be of two minds about this.

Here is what White himself says to justify modeling the case this way. His idea is that he's focusing on subjects who not only have handlike experiences, but are also reflectively aware of having them. These subjects are better informed about their own epistemic situation. Surely that should not make them *worse* off with respect to whether they have hands, right? True, philosophers have sometimes told stories where knowing more can make you epistemically worse off overall; but there's no apparent reason to think these cases are like that. So if we find constraints on how justified even these better-informed subjects can be, we should expect those constraints to apply with at least much force to less-informed subjects, too.³⁸

White's strategy will be to identify such a constraint.

In fact, given the specific arguments he wants to present, White's hands were sort-of tied to proceed in this way. His arguments posit a skeptical hypothesis U that both *entails* the proposition you update on, but is *incompatible with* the proposition it undermines. That posit could never succeed in a case where what gets undermined is the same proposition you update on. So White's presentation seems to *require* modeling the dogmatist's idea that your immediate justification is underminable in a way that *doesn't* involve you updating on the proposition you're acquiring immediate justification for.³⁹

We've mentioned and set aside three of the threads in White's discussion. A fourth thread is what we will focus on. This thread can also be presented using the argument from E to Good; but it targets all dogmatists about perception, regardless of their attitude towards such "Moorean" reasoning.

³⁸See White 2006, pp. 534–5; see also section 3 of Wedgwood forthcoming; and Silins 2008, note 22. Silins raises some worries in his forthcoming, note 30 and preceding text.

I've been disposed to go along with White here; his strategy seemed to me a reasonable one. But recently David Barnett has persuaded me that the issues here are not so straightforward. White's proposal encourages the idea that the subjects are just coming to discern facts about the epistemic situation that were already in place. It's as though we asked them "How many US States have names beginning with the letter M?" The relevant information is already there in their mind; reflection just has to locate it. Contrast a case where we ask the subject "How many US States remind you of your grandmother?" Here too, they may be able to answer the question without leaving the armchair. But most subjects would perform mental experiments to answer the question, or engage in other mental activity that intuitively *changes* their epistemic situation as they proceed. Are the better-informed subjects White focuses on more like the first group of subjects? or more like the second? That's not entirely clear. And if they're more like the second group, then the differences between them and less-informed subjects might not be so harmless.

³⁹As Maria Lasonen-Aarnio pointed out to me, though, White only *uses* U's entailing the proposition you update on in order to secure the result that $\text{Old}(U|E) > \text{Old}(U)$. He might substitute a different U that retains that property, but doesn't entail the proposition you update on.

This criticism turns on the claim that your new probability for Hand—the credence you’re justified in having after updating on E—cannot be higher than your prior probability for Good. The proof of this is not hard:

Old(.) is your epistemic probability function before learning E or Hand. New(.) is your epistemic probability function after E becomes true. White has defended his choice of letting New(.) be the result of conditionalizing on E; that is, $\text{New}(\cdot) = \text{Old}(\cdot|E)$. We note that observations that are entailed by a hypothesis, as E is entailed by Bad, contribute positively to the probability of that hypothesis. That is, $\text{Old}(\text{Bad}|E)$ will be greater than $\text{Old}(\text{Bad})$ when E is not yet itself epistemically certain. This is equivalent to:

$$\text{Old}(\text{not-Bad}|E) < \text{Old}(\text{not-Bad}).$$

Next we observe that, since you have updated on E:

$$\text{New}(\text{not-Bad}) = \text{Old}(\text{not-Bad}|E).$$

And further, we observe that Bad is incompatible with Hand. In other words, Hand entails not-Bad and its probability can be no higher than not-Bad’s:

$$\text{New}(\text{Hand}) \leq \text{New}(\text{not-Bad}).$$

Putting these three equations together, we have:

$$\text{New}(\text{Hand}) \leq \text{New}(\text{not-Bad}) = \text{Old}(\text{not-Bad}|E) < \text{Old}(\text{not-Bad}).$$

Ignoring the middle terms, we have:

$$\text{New}(\text{Hand}) < \text{Old}(\text{not-Bad}).$$

So Bayesianism tells us that your new probability for Hand cannot be higher than your prior probability for not-Bad, that is, your prior probability for Good.

Though White doesn’t argue it, an even stronger formal result is possible. We can show, not merely that $\text{New}(\text{Hand}) < \text{Old}(\text{not-Bad})$, but also that $\text{New}(\text{Hand}) < \text{Old}(E \supset \text{Hand})$, where the latter quantity may be even lower than $\text{Old}(\text{not-Bad})$. So this is a constraint that will bind at least as tightly as the original, and sometimes more tightly.⁴⁰

⁴⁰Proof: $\text{Old}(E \supset \text{Hand}) = \text{Old}(\text{not-E}) + \text{Old}(E \wedge \text{Hand}) = \text{Old}(\text{not-E}) + \text{Old}(\text{Hand}|E)\text{Old}(E) = \text{Old}(\text{not-E}) + \text{Old}(\text{Hand}|E)(1 - \text{Old}(\text{not-E})) = \text{Old}(\text{not-E})(1 - \text{Old}(\text{Hand}|E)) + \text{Old}(\text{Hand}|E) = \text{Old}(\text{not-E})\text{Old}(\text{not-Hand}|E) + \text{Old}(\text{Hand}|E)$. When $\text{Old}(\text{not-E})$ and $\text{Old}(\text{not-Hand}|E)$ are each > 0 , this will be $> \text{Old}(\text{Hand}|E)$, which is $\text{New}(\text{Hand})$. So in those circumstances, $\text{New}(\text{Hand}) < \text{Old}(E \supset \text{Hand})$. Finally, since Bad entails that E and not-Hand, $\text{Old}(E \supset \text{Hand})$

One can see from the way these results were derived that nothing here is specific to perception. Any case where one thinks there is underminable immediate justification should display the same structure. For example, perhaps I have introspective justification to believe I intend to confront my father about something. Then I read a psychological study that says subjects are unreliable about whether they even intend, as opposed to merely fantasize, such things. I presume that would undermine my introspective justification. And reasoning just like White's would seem to establish the same inequalities. Arguably even my justification to believe E would also have this structure.

Well, what is the significance of the inequalities we've established?

Here is how White interprets them:⁴¹

So its appearing to me that this is a hand can render me justifiably confident that it is a hand, only if I am already [justifiably] confident that it is not a fake-hand. (p. 534)

In taking this to tell against dogmatism, White commits himself to what I'll call:

Assumption-4 If you *already* need to be justified to a certain degree in believing A, in order to acquire some new quantity of justification to believe B, then it's false to say that your new justification for believing B is immediate. It did after all need to include at least that degree of antecedent justification to believe A.

To sloganize it, we might think of this as the idea that **prior probabilities have epistemic antecedence**.⁴²

\leq Old(not-Bad).

It is more difficult to prove in the Jeffrey framework that $\text{New}(\text{Hand}) < \text{Old}(E \supset \text{Hand}) \leq \text{Old}(\text{not-Bad})$, however the same result does also hold there.

Bad in these examples is generally assumed to entail that I have the specific experiences I do (though see note 39). Presumably the prior probability of me having specifically experiences E will be extremely low; so not-Bad's prior probability will be extremely high. As David Christensen reminded me, though, it also follows that the prior probability of $E \supset$ anything will be extremely high. So $\text{Old}(E \supset \text{Hand})$ shouldn't be expected to usually be *much* lower than $\text{Old}(\text{not-Bad})$.

⁴¹White clearly intends the inserted "justifiably." If we genuinely do consider the effects of mere, possibly unjustified, confidence that Good, this introduces new issues. See the papers cited in note 32 for discussion.

⁴²Assumption-4 makes two steps, one signaled by the word "include" and the other by the word "antecedent." The first of these steps was mentioned in note 17. Silins 2008, section 4, argues that this step is not mandatory, and I agree. The second step takes the justification represented by prior probabilities to be "antecedent" in the sense we're working with—regardless of whether it's an included part or merely a necessary condition of the justification represented by your posterior probabilities. This is what I will focus on.

Willenken 2011 calls Assumption-4 "ATIJ" and argues against one variation of it. The notion of justification I'm working with is sensitive to his "reasoning-directed defeaters," and

This idea promises to be as problematic for credulism more generally as it is for dogmatism in particular. Any attempt to model non-quotidian undermining in the Bayesian framework is going to face structural choices like the ones we've been discussing. It's going to be natural to (try to) model cases with non-quotidian underminers in ways that don't look any *formally* different from how we model quotidian undermining. For example, suppose you believe H on the basis of E, and "E seems not to support H" is a non-quotidian underminer. One expects a Bayesian model of this to be such that:

$$\text{Prob}(H) \leq \text{Prob}(H|E \wedge E \text{ seems not to support } H) < \text{Prob}(H|E)$$

In just the same way that:

$$\text{Prob}(\text{It will rain}) \leq \text{Prob}(\text{It will rain}|\text{This barometer is falling but is unreliable}) < \text{Prob}(\text{It will rain}|\text{This barometer is falling})$$

If we then subscribe to Assumption-4, we'll interpret these structural facts as implying that, in fact, you need justification for "E seems to support H" epistemically antecedent to any justification for H. So this case turns out to be not one of non-quotidian undermining after all.

Surely Assumption-4 is philosophically optional. If it's going to be possible to model non-quotidian undermining in Bayesian terms, we must do so without that Assumption in place. So why should a dogmatist or credulist agree to it?

And we can do better than complain that Assumption-4 is optional. We can argue that White has himself already presumptively relied on its falsity; and we can show that in the general form we've stated it, it has intolerable consequences.

so is what he engages with under the heading of "robust" views. However, because of the issues discussed in the papers cited in note 32, I officially have what he defines as an "anemic" view of justification.

Some care is needed in attributing Assumption-4 to White. What's clear is that: (i) his paper nowhere attends to, and is sometimes insensitive to, the difference between its being a necessary condition that you have some prior probability, and your justification including antecedent justification to believe something else. It's also clear that: (ii) the theories he's taking his formal result to arbitrate between are defined by their proponents in terms of the second notion, so something like Assumption-4 is needed to make the formal result relevant. After that, things are less clear: (iii) the language in which White himself defines the theories he's considering can in places be read as invoking the second notions; but it can also be read just in terms of the first notions, and White's paper strongly suggests that this is his intent. (See esp. his note 14; and he has confirmed this in discussion.) So the best way to understand him may be as arguing for a position like we'll entertain in the next section. He just incorrectly takes that position to be incompatible with dogmatism, and to suffice for the kind of view espoused in Wright 2002.

For two reasons, I won't keep re-acknowledging this. Instead I will proceed as though White's criticisms are intended to oppose dogmatism as I understand it, and so he really does commit himself to Assumption-4. The first reason is that many of his readers have understood him that way, and sometimes endorsed what they so understood. The second reason is to keep our dialectic more manageable. It is instructive to figure out how the criticism I'll treat White as advancing fares. The biographical facts about what he or I ever intended are of less interest.

In what sense has White “presumptively relied” on Assumption-4 being false?

I invite you to compare it to Assumption-2, which said that the immediately justified propositions are what you update on. These two ideas seem to naturally go together. Yet White rejects Assumption-2—and as we saw, he *needs* to do so in order for his model of the situation to be charitable to the dogmatist. Remember: if all the business about probabilities is to do any work—if there’s to be more to his argument than the bare assertion that when you have handlike experiences, you don’t get immediate justification to believe Hand—then the mere fact that Hand goes up as a function of your updating on something else shouldn’t be understood to mean that the justification you acquired for Hand was epistemically posterior to anything else. So what White has relied on is that *the functional dependencies in the formalism can’t in general be assumed to mirror the facts about epistemic dependence and priority*. Assumption-2 and Assumption-4 turn on different functional dependencies. But if the dependencies of the latter justify an interpretation that those of the first do not, this needs special motivation.

And when you think about it, it doesn’t seem like Assumption-4 can be true, not in the general form we’ve stated it. The implications are much too strong. Let H be any hypothesis whose probability goes up when you learn that E . Let H^* be any old logical implication of H . Now it’s straightforward that:

$$\text{New}(H) \leq \text{New}(H^*) = \text{Old}(H^*|E)$$

Moreover, whenever $\text{Old}(E)$ and $\text{Old}(H^*|E)$ are both less than 1, then $\text{Old}(E \supset H^*)$ will be strictly greater than $\text{Old}(H^*|E)$.⁴³ So we can add:

$$\text{New}(H) \leq \text{New}(H^*) = \text{Old}(H^*|E) < \text{Old}(E \supset H^*)$$

That is, you can’t acquire justification to believe H above a given threshold, in response to some evidence E , unless you already had at least that high a prior probability that $E \supset H^*$ —for *any* implication H^* of H . Do we really want to conclude that the justification you acquired for H was partly constituted by or came from antecedent justification to believe *each* such conditional?⁴⁴

Rejecting that interpretation of the inequality $\text{New}(H) < \text{Old}(E \supset H^*)$ means rejecting Assumption-4. Perhaps some more restricted version can fare better, but we’ll have to wait and assess it when we see it.

Without some such Assumption, the formal results we derived don’t directly threaten anything a dogmatist or credulist wants to say. In the next section, I’ll spell out some ways of thinking that develop the idea of prior probabilities not having epistemic antecedence in the way that Assumption-4 envisages.

⁴³See note 40.

⁴⁴Thanks to Philip Ebert for stressing the awkwardness of this to me.

7. Life without Assumption-4

What might explain an inequality like $\text{New}(H) < \text{Old}(E \supset H^*)$, or $\text{New}(\text{Hand}) < \text{Old}(\text{not-Bad})$, if not the fact that the justification you acquire for the former proposition is partly constituted by the antecedent justification you have for the latter?

Consider what Boghossian 2000 says about our epistemic relation to modus ponens. We may well be justified a priori in believing:

Valid. If P, and $P \supset Q$, are true, then Q must be true.

However our justification for applying the rule of modus ponens in a deduction doesn't *come from* our justification to believe Valid. Rather our justification to believe Valid relies instead on the reasonability of our applying the rule. There's an important question of epistemic priority here, that isn't settled by the fact that our justification for Valid is a priori. (See also Boghossian 2001 and 2003.)

We don't need to assume that Boghossian's story about that is right; nor that, even if it is right, the same story extends to our present concerns. I merely ask you to attend to the *epistemic structure* he posits between our justification to believe Valid and our justification to reason by modus ponens. A similar structure could obtain in the case of perception, or anywhere else we have immediate justification. That is, even if belief in $E \supset H$ is a priori, our justification for believing H in response to E needn't *come from* our justification to believe that conditional. The epistemic priority may in fact be reverse.

Several philosophers have thought that wherever a transition from some evidence E to H is epistemically legitimate, the conditional $E \supset H$ *must* be a priori justifiable.⁴⁵ A dogmatist needn't oppose this. Whether it conflicts with anything

⁴⁵See BonJour 1998, section 7.7; and Hawthorne 2002. This is an idea with a longer history, though I'm not sure how much longer. Chisholm held that something of the form "If . . . then $E \supset I$ am justified in believing H" was synthetic a priori (see his 1989, at pp. 72–3. See also van Fraassen 1989, section 6.3. I heard the epistemic descent from "I am justified in believing H" to H made several times in the 1990s.

The formalisms we're working with require a proposition to have some defined initial probability if it's ever to have any posterior probability. Philosophers who interpret the probabilities as representing justification, as we are, often interpret the initial probabilities as representing what you have a priori justification to believe. I won't resist this; though I will point out a different interpretation of "a priori" below. But in fact I have serious doubts about this practice. The formalism requires initial probabilities for E, H, and $E \supset H$. Am I sure that I had a *a priori justification* to have any particular confidence in these claims? for example, the claim E, that I'd have just these experiences now? This is an extraordinary claim.

Perhaps there may be ways to understand the probabilities as representing facts about justification, without interpreting the initial probabilities that way. If Assumption-4 is optional, this may well be optional too. (Kung 2010, at p. 12, also suggests not interpreting a high prior in Good to mean you have any justification to believe Good.)

If it's not optional, then we should at least want to move to versions of these formalisms that can represent probabilities as initially very imprecise; even better would be ones that permit probabilities to be initially undefined. Claims like Assumption-4 would then be even less straightforward.

he says depends on whether it's the justification of the conditional that underwrites the reasonableness of the transition. It may be the other way around. (Or maybe neither underwrites the other.)

Similarly, a credulist needn't oppose its being a priori that $(E \wedge E \text{ supports } H) \supset H$. That only forces him to say his E-based justification for H includes antecedent justification to believe "E supports H" if his justification for H is *epistemically posterior* to his justification for this conditional. It need not be.

Just because one claim is a priori and another isn't, it doesn't follow that the second is epistemically posterior to the first—that your justification for it needs to include antecedent justification for the first. Not even when the claims' contents are relevant to each other. There may *be* routes to the second that proceed by way of the first; but there may also be routes that don't.

The dogmatists and credulists think that the reasonableness of your move to H doesn't derive from antecedent justification to believe the conditionals we mentioned. When you learn E, you don't need to rely, not even implicitly, on an application of modus ponens. Doing so would be a mistake akin to Achilles' mistake in agreeing that Valid is another premise he needs to argue that Q.

Cohen and Wedgwood have recently argued for the a priority of $E \supset H$ as a result of using non-deductively good reasoning inside a conditional proof.⁴⁶ Their idea is that if it's prima facie reasonable to conclude you have hands when you *really do* have experiences as of hands, then it's also prima facie reasonable to conclude you have hands *under the supposition that* you have experiences as of hands. And then you can discharge the supposition to infer that:

You have experiences as of hands \supset you do have hands.

It's not obvious that those epistemic transitions are legitimate. The epistemic effects of having some experiences needn't be the same as, or even inherited by, the proposition that you have them. Or there may be general reasons why we can't expect arbitrary good reasoning moves to also be licensed in suppositional contexts. Certainly some reasoning moves that have been claimed to be good are not so licensed. Alex Byrne claims that the transition from P to I believe P is justificatory (Byrne 2005 and 2008). Perhaps he's wrong, or perhaps not. We should agree, though, that under the supposition P I should *not* be entitled to conclude that I believe P, and then use a form of conditional proof to infer that $P \supset I \text{ believe } P$. As David Barnett points out, if I could do that I should also be entitled to argue in the same way that $\text{not-}P \supset I \text{ believe not-}P$. For any P at all. And the conjunction of all those conditionals looks like the claim that I am omniscient.⁴⁷

⁴⁶See Cohen 2010 and Wedgwood forthcoming. White 2006 section 6 anticipates their proposals in some ways.

⁴⁷Wedgwood acknowledges this specific limitation; see his note 13.

Moreover, Weatherson 2012 gives a compelling case where a very weak inductive rule looks to be illegitimate in a suppositional context.

So it's not clear to me that the Cohen/Wedgwood strategy is correct. But some other story about why we're a priori justified in believing $E \supset H$ may be. Such stories need not make the reasonableness of moving from E to H epistemically dependent on our justification for the conditional. The story might make that justification parasitic, in some other way, on the reasonableness of moving directly from E to H. (Or maybe neither is epistemically posterior to the other.)

Juan Comesaña discusses the possibility of combining such stories with a commitment to immediate justification.⁴⁸ He says:

[T]his position is not very stable... [W]henver you are justified in believing H on the basis of E you will have available to you a stronger justification---one which depends not only on E but also on the conditional if E then H.

Even if the roundabout, deductively valid argument for H were available whenever the immediate justification is, why should we count the roundabout route as a "stronger" one? Comesaña complains:

Of course, the friend of immediate justification is free to hold that even though you have justification for believing this conditional, your justification for believing the consequent 'comes from' just the antecedent. ...What determines whether your [propositional] justification for believing H comes from (in part) the conditional or not? We could try defending the claim that it doesn't by saying that you would still be justified in believing H even if you were not justified in believing in the conditional. But remember that, according to the neo-Rationalist, justification for the conditional will be available whenever the antecedent justifies the consequent...

But it seems readily imaginable for you to have the immediate justification but lack justification for the conditional: they may have different defeasibility profiles. Let your philosophical mentor give you a mistaken but compelling argument against the conditional, which allegedly doesn't threaten the immediate justification. As we said in section 4, what effect such evidence should have on your other justification is controversial. But a natural view is that it would drag the justification you have for the conditional in a defeated direction, and need not defeat the immediate justification you have from E to the same degree.⁴⁹

So the combination of belief in the a priority of $E \supset H$, and continued commitment to immediate justification, looks entirely feasible to me. The a priority of some

⁴⁸Comesaña forthcoming; I've changed the labeling in the quotations.

⁴⁹See also my 2000, note 6.

such conditional doesn't establish its epistemic antecedence over anything. At best, it only establishes its *lack of epistemic posteriority* on certain kinds of evidence.

Our notion of “a priority” may have an indeterminacy that blinds us here. Since the 1980s, it has come to be widely appreciated that some experiences may be necessary for the having of a belief, without threatening that belief's claim to be a priori.⁵⁰ For example, maybe no one can have the concept of *umami* who's never had taste experiences. But the beliefs that umami things are umami, or that integers are not umami, are presumably a priori. Experiences are necessary for the having of these beliefs, but don't contribute to their justification. (Other times, experiences that are necessary for the having of beliefs *can* contribute to their justification.) What the preceding reflections suggest is a converse possibility, that experiences may play a role in the justification of some beliefs even without ever being had. You may be justified in believing $E \supset H$ in part because of the justificatory power of experiences E—even though you don't now have E. Perhaps you can be so justified without *ever* having had E, but just by reflecting on what it'd be epistemically like *were you* to have E. Does that make your justification for the conditional $E \supset H$ a priori?

I feel pulled in two directions here, and that suggests that my concept of a priority can be extended in either direction. We might say: justification is “a priori” if it can be had *in advance of* having any experiences.⁵¹ So understood, our justification to believe $E \supset H$ would be a priori, because it can be had in advance of ever having experiences E. Despite that, though, it may still be the epistemic powers of E that justify the conditional, rather than your justification for the conditional justifying you in concluding H when you really do come to have E.

Alternatively, we could use “a priority” to track whether the epistemic powers of any experiences or experiential capacities contribute to the justification of your belief. Understood in that way, your justification for these conditionals is not a priori, even if it could be in place before you've ever had any experiences, and so even if it could be reflected in your initial probability function. It's not a priori because it's *parasitic on* what the experiences can justify you in believing, when you do have them.⁵²

In sum, I think it's open to discussion whether prior probabilities in the Bayesian sense should be understood *either* as a priori *or* as having any epistemic an-

⁵⁰Kant also emphasized this; but the point was long unappreciated in the 20th century.

⁵¹Though we shouldn't also make that a *necessary* condition, because we still want “Integers are not umami” to be justified a priori, and maybe I can't have that justification until I have the concept of umami, which may require having experiences.

⁵²Yablo 2002, sections 12 and 15, argues that imaginative judgments about ovals fail to be a priori even though they don't evidentially rely on any sense-experiences. (Neither does it seem right to say they evidentially rely on the deliverances of introspection.) Williamson 2007, pp. 165 ff. argues similarly about imaginative comparisons of inches to centimeters. I understand them to be working with something in the vicinity of this second notion of “a priority”—though we may also want to make finer-grained distinctions.

tecedence. White and others have been assuming that the debate between dogmatists and their opponents boils down to what our prior probabilities have to be. Whether it does so or not depends on how we settle some of these open questions.

It may for all I know be the best course for a dogmatist, to agree with White about the prior probabilities, but to wrestle about what their philosophical significance is. This may also be an open avenue for credulists who aren't dogmatists. They may agree to attribute high prior probability to claims like "E supports H," and "I am a Competent Reasoner," without your justification for those claims yet playing a constitutive role in the justification of the inferences you make.

8. Jeffrey Bayesianism

We've looked at some philosophical assumptions behind the Bayesian-based objections to dogmatism. We've seen that these assumptions are not mandatory. However, neither is it really straightforward how to model the possibility of immediate yet underminable justification—especially if we retain Assumption-2, and so think that that needs to involve an underminable *update*. We only gestured at some ways this might be understood. As we said before, some of the ideas we were gesturing at look like they'd be more promisingly developed in a Jeffrey Bayesian framework. So let's turn our attention there.

This section will say a bit to explain how this framework works. The next section will summarize some arguments due to David Christensen and Jonathan Weisberg that suggest that Jeffrey Bayesianism *isn't* after all very hospitable to the possibility of underminable updates. These arguments also involve some substantial philosophical assumptions, which I will identify.

OK, so what is Jeffrey Bayesianism?

Suppose you start off with a probability distribution that looks like figure 2.

Now in the Classical framework, if you update on E, then your probability distribution will change as in figure 3.

That is, we just erase the not-E part of your old distribution and retain the E part. All the relative relationships *within* the E part remain the same. That is, when you update on E, then for any proposition H:

$$\text{New}(H|E) = \text{Old}(H|E)$$

This fact is called **your update being rigid with respect to E**.⁵³

⁵³In the Classical framework, the left-hand term of the equation is equivalent to $\text{New}(H)$ simpliciter, but that won't also be true in the Jeffrey framework.

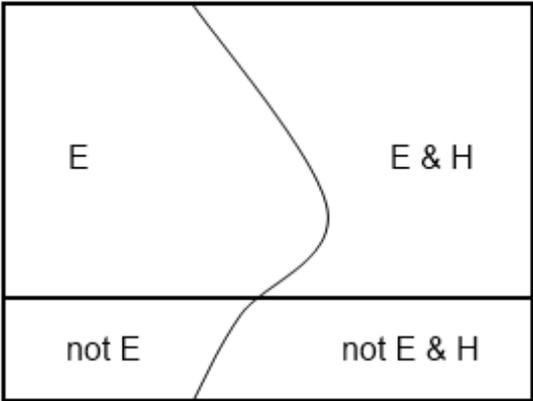


Figure 2:

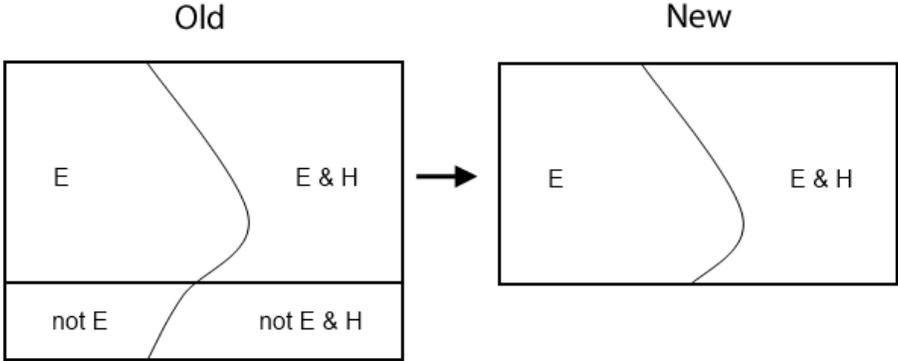


Figure 3:

Now Jeffrey Bayesianism shares much of what we just described, and has the Classical update as a limiting case. Jeffrey's transitions look as in figure 4.

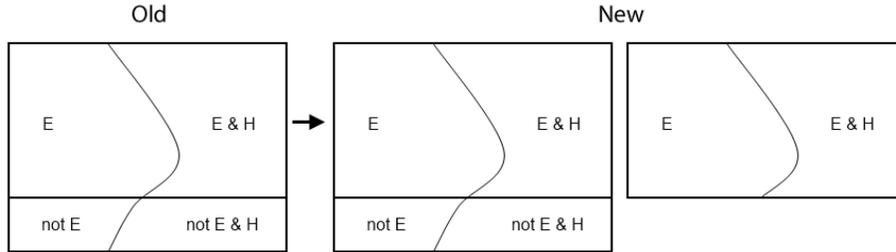


Figure 4:

That is, we don't just get a copy of the E-part of our old distribution. We also get a copy of the entire old distribution too. Of course, there will be a question of the relative sizes of these two components. That might not always be the same fixed ratio. When specifying an update, it somehow needs to be settled *which proposition* you update on (in this case, E), and also *what scaling factor should be applied* to the right-hand component of your new distribution. Let's suppose the two components are scaled as in figure 5.

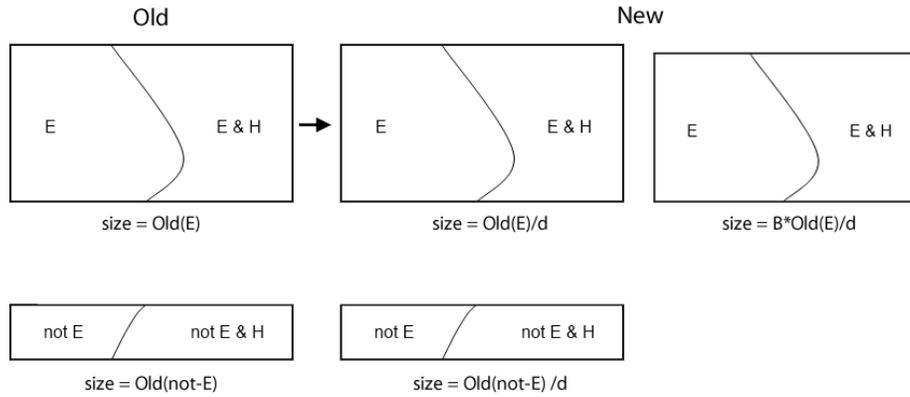


Figure 5:

Where B is some real value, and d is a normalizing factor so that $\text{Old}(E)/d + \text{Old}(\text{not-}E)/d + B*\text{Old}(E)/d = 1$. In figure 5, B looks to be roughly 1. But we can suppose it to be any non-negative real. As B approaches 0, New's right-hand component becomes relatively much smaller, and it's as though you didn't become much more confident of E at all. As B approaches infinity, the right-hand size dwarfs the left-hand, and we approach the situation we had in the Classical case, where you moved just to a copy of the E-part of your original distribution.

Perhaps there is some magic single value for B that all updates should use. Or perhaps, as most theorists assume, different learning episodes would make different B s appropriate. A good look at an object in noon sun may call for updating on the proposition that it's red, with a high B ; a quick look in twilight may call for doing so with a low B .

It's part of the Jeffrey machinery that these transitions are also rigid with respect to E . The two left-hand components of your new distribution are internally just the same as they were in your old distribution; and the right-hand component retains all the internal relationships of $\text{Old}(\cdot|E)$, just as we saw in the Classical update. So in the Jeffrey framework, it's also true that an update on E is rigid with respect to E .

This story is simplified in some ways; but it is the essence of what many now understand by "Jeffrey updating."⁵⁴ What stands out is that when you update on a proposition, as we did here on E , you needn't become certain of that proposition, but merely more confident of it. That leaves open the possibility that you might later acquire other evidence that opposes that proposition. Perhaps it might also leave open the possibility that your update on E could be defeated by being undermined? Let's wait and see.

An interesting question is what it would be for someone who started with a different probability distribution to update *in the same way* as we just illustrated. It's natural to expect this will involve also updating on the same proposition E . But how big should the scaling factor be, for it to be the same update?

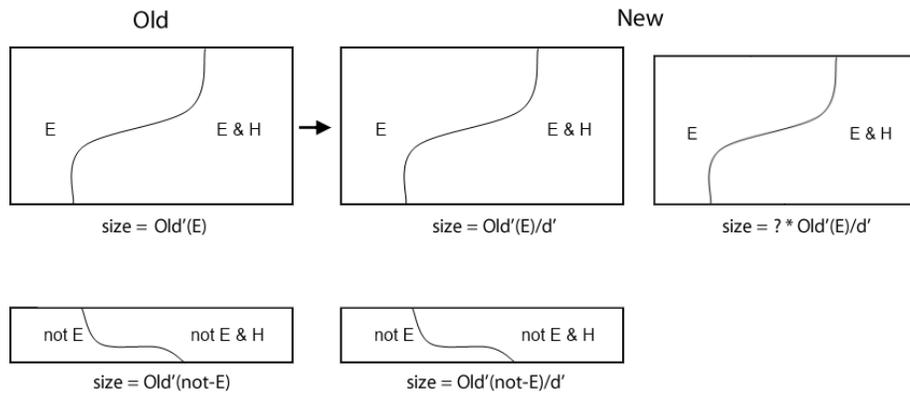


Figure 6: What does $?$ need to be, for this to constitute updating *in the same way* as before, albeit from the starting point of $\text{Old}'(\cdot)$ rather than $\text{Old}(\cdot)$?

It turns out that, if we want updates to **commute**, it is necessary and sufficient for us to have *the same* update that it be an update *on* the same proposition and *using the same scaling factor* B . This may look natural in the diagrams we

⁵⁴The model presented here purports to be exhaustive in a way that Jeffrey would not himself endorse; see note 60.

have, but that's because I deliberately explained Jeffrey updating in a way that makes it look so. The result wasn't at all initially obvious.⁵⁵

But should we want evidence to commute? Isn't it sometimes significant which of two pieces of evidence we acquired first? For example, what if we knew in advance that our measuring devices deteriorate rapidly? Then the first use of them is more likely to be reliable.

Yes, of course. But that worry is not envisaging two scenarios where we have all the same evidence, acquired in different order. It's imagining that in one scenario, you have the evidence that *you got instrument reading A first*, and in the other scenario, evidence that *you got instrument reading B first*. So those are cases where your total evidence is different. Who would deny that such a difference may sometimes be epistemically relevant? What would be needed for a case that demonstrated an intuitively compelling violation of commutativity would be a case where in each scenario you had all the same evidence, including the same evidence about which evidence you acquired first, but *the facts* about which evidence you acquired first were different, and *that* by itself made an intuitive difference to what it was reasonable for you to believe. It's not clear there could be any such intuitively persuasive case.⁵⁶

I think the best way to think about commutativity is to think about its application in cases where *there is no* empirically significant difference in the order in which you acquire two pieces of evidence. Instead, the choice of which to update on first is just a formal artifact. For example, you may have your left hand and your right hand in each of two holes in the ground, and you simultaneously feel something with each hand. Shouldn't we get the same results by doing any of the following:

- first update on your right-hand evidence, then on your left-hand evidence, then on the claim that they were acquired simultaneously
- first update on your left-hand evidence, then on your right-hand evidence, then on the claim that they were acquired simultaneously
- update in the first place only on some composite of your left- and right-hand evidence

At root, I think commutativity just amounts to a commitment that our formalism will not attach any significance to which of these routes we take. Whether

⁵⁵For a time, it was widely accepted that Jeffrey updating must be non-commutative. Field 1978 demonstrated that making the update partition and B invariant suffices for commutativity. Wagner 2002 extended Field's result, and showed that on minimal assumptions this invariance is also necessary for commutativity. See also Hawthorne 2004a.

⁵⁶For externalists about evidence, it may prove difficult to even *describe* such a case, since they allow that such facts could make a difference to what evidence you have.

Lange 2000 criticizes putative demonstrations of (Jeffrey's commitment to) non-commutativity in roughly the same way I do in the text.

empirically real differences in the order we acquire different pieces of evidence makes an epistemic difference, on the other hand, depends upon the details of our situation. That issue shouldn't be confused with commutativity.

So now we've gotten an intuitive grip on what Jeffrey Bayesianism involves; we've observed that Jeffrey updates, like Classical updates, are rigid with respect to the proposition being updated on; and we've heard that commutativity entails that updating "in the same way" means updating with the same scaling factor B .

That's all the machinery we need, to see that Jeffrey Bayesianism ultimately isn't that hospitable to the key ideas of dogmatism, either.

9. Threats from Jeffrey Bayesianism

There are three arguments for this conclusion, using somewhat different, plausible but still controversial assumptions.

The first of these arguments is from Weisberg 2009, and turns on facts about rigidity and independence.⁵⁷ Recall, rigidity is the property that when you update on E , all your conditional probabilities given E , and given not- E , remain the same. This is a core feature of Classical Bayesianism, of Jeffrey Bayesianism, and of some other popular formal models of rational credence as well. It's sometimes motivated by saying:

Assumption-5 When what you learn is E , you don't learn anything about the epistemic relations between E and other propositions.

However, whether that intuitive Assumption really is correct, or whether it really is equivalent to the formal property we've identified, seem substantial questions. I will not pursue them.

Weisberg presents the problem this poses as follows. We suppose you will undergo some learning episode that will involve you Jeffrey updating to some degree on E , but this hasn't yet happened. We suppose also there could be a *pure underminer* U for the support you thereby acquire to believe E . This is a notion we discussed back in section 2: an underminer not mixed with any opposing or supporting elements.⁵⁸ It's at least initially natural to think this should be possible. It's also natural to think that what this amounts to is that, where $\text{Old}(\cdot)$ is your probability function before the learning episode:

⁵⁷Weisberg's paper is titled "Commutativity or Holism?..." "Holism" here is roughly the thesis that all acquisitions of (at least empirical) justification are underminable. Weisberg forthcoming extends his arguments to some other formal models of belief update.

⁵⁸Hence, U in Weisberg's argument is very different from U in White's argument. White's U was *incompatible with* the proposition H and whose justification we were considering.

$$\text{Old}(E|U) = \text{Old}(E)$$

That is, before the learning episode takes place, we expect U to be probabilistically independent of E. Getting more evidence for U should not yet either increase or decrease your credence in E. Let's label this:

Assumption-6 If the support you acquire for updating on E is underminable at all, then it could be undermined by some U such that $\text{Old}(E|U) = \text{Old}(E)$.

Now let the learning episode take place. It is the very proposition E you're updating on whose support U undermines. If we keep Assumption-2's identification of what you should update on with what you acquire immediate justification for, then this is just the kind of possibility the dogmatist posits. Let $\text{New}(\cdot)$ be your probability function after the learning episode. Observe that we *wouldn't* think that U should be probabilistically independent of E *in* $\text{New}(\cdot)$. Instead, the higher your credence was in U, the *less* support we should think you end up with for E in $\text{New}(\cdot)$. If U *thoroughly* undermines the support you newly acquired for E, then $\text{New}(E|U)$ should leave you back at $\text{Old}(E)$. But it is not essential for U's undermining effect to be so thorough. All that is important is that we think, now:

$$\text{New}(E|U) < \text{New}(E)$$

However, what this means is that U and E started out independent, and then ended up no longer independent, as a result of a Jeffrey update on E, which by definition will be rigid with respect to E. And what Weisberg observes is that that is impossible. If an update is rigid with respect to E, then whatever started out independent of E must end up independent of E.

So if Assumption-2 and Assumption-6 are right, then the possibility the dogmatist posits is not representable in Jeffrey Bayesianism, after all. You can represent the support you acquire for E, when updating on E, as being *defeasible*—for example, you could go on to acquire *opposing* evidence for not-E. But you can't represent this support as being purely underminable in the way envisaged by Assumption-6.

There are ways out of this. We might decide that pure underminers aren't possible after all, or that Assumption-6 doesn't correctly formally capture what they involve. Or we might try one of the strategies from section 5, and give up Assumption-2 and/or -3. What you gain immediate justification for might not be the same as what you update on. Immediate justification might not correspond to *any* specific update. The negative effect of undermining evidence might not be wholly represented *inside* the formal model: it might help determine *what* you update on, or to what degree. Maybe the appropriate scaling

factor should be a function of how much justification you independently have for U.⁵⁹

But something needs to give.⁶⁰

I said there were three arguments. The second and third arguments turn not on the requirement of rigidity but rather on the requirement of commutativity.⁶¹

Theorists argue about what is the best way to measure the amount of evidential support some update contributes to a given proposition E. One natural such measure is called the *E:not-E Bayes factor* of the update. This is not the only measure that has advocates, but it does have some broad support. We needn't dwell on the details of how this measure is defined.⁶² The bottom line is that the E:not-E Bayes factor of an update on E will turn out equivalent to $1+B$, where B is the scaling factor displayed in our earlier diagrams. That's the same scaling factor we said the Jeffrey formalism *forces to be invariant*, for subjects who update "in the same way," on pain of violating commutativity. So in other words, if you and I have a good look at an object in noon sun, and it looks the same to us—and that means that our credal models should be supplied the same input—then we'll need to update on some single E and do so with the same E:not-E Bayes factor; though our prior probability functions will differ. But then if Bayes factors are also the right way to measure how much evidential support our updates contribute to E, it will follow that each of us will have acquired *the same amount* of such support. But we've said nothing about what other ways we might differ! Old'(U) might much higher for some

⁵⁹I have developed an extension of the Jeffrey formalism which works that way.

⁶⁰Wagner forthcoming criticizes Weisberg's argument, making the following points: (i) Jeffrey did not intend the updating rule we described in section 8 to be the "be-all-and-end-all of probability revision." Moreover, (ii) you can reasonably update with that rule on a proposition E only when it is *antecedently* reasonable for you to believe that the conditional probability for any U on E should be unaffected by what you've learned. Finally, (iii) in cases like the ones we're discussing, what it's reasonable for you to update on is not E but $E \vee U$.

I'm prepared to believe that (ii) was Jeffrey's view. Certainly it is a natural way to read the end of Jeffrey 1987. But if point (ii) is correct, it severely limits the applicability of the model described in section 8. I suspect that only theorists who embraced the "towers of antecedent justification" we described in section 4, or theorists who thought we only acquired non-underminable justification—that is, only non-credulists—would then be able to regard the model as an exhaustive theory of reasonable belief-update. (As Wagner points out, Jeffrey did not himself so regard it.)

Wagner's point (iii) tacitly concedes the main point I'm taking Weisberg to be arguing: this model is not suited to represent updates where it's your justification for the very proposition E you're updating on that's underminable. (Wagner instead recommends updating on a proposition whose acquired justification *isn't* underminable by any of the other hypotheses in your algebra.) So long as our Assumption-2 is in place, this means that the model is not suited to represent updates of the sort dogmatists believe possible.

⁶¹I've been helped in thinking about the second argument by discussions with Matt Kotzen. The third argument comes from Christensen 1992, section 6, and is also presented in Weisberg 2009. Christensen's paper is titled "Confirmational Holism and Bayesian Epistemology." "Holism" here can be understood in the way explained in note 57.

⁶²For reference, it is the quantity $(\text{New}(E)/\text{New}(\text{not-}E))/(\text{Old}(E)/\text{Old}(\text{not-}E))$. Sometimes the log of this quantity is used instead. In a Classical setting, this is equivalent to the "likelihood-ratio" measure of evidential support.

underminer U than $\text{Old}(U)$: that is, I may have much more evidence than you do for thinking vision is unreliable in noon sun. Shouldn't that have *undermined*, and so weakened, the support I get for E, when my credal system gets the same input as yours?

We have a conflict between the following ideas:

- (i) The dogmatist's claim that what you acquire immediate justification for is itself underminable.
- (ii) Assumption-2's claim that a proposition you acquire immediate justification for is what you update on.
- (iii) If the support I acquired for E is undermined to a greater degree than the support you acquired, then the measure of *how much* support I got for E should be less.
- (iv) The right way to measure that support is in terms of the E:not-E Bayes factor of the update.
- (v) But commutativity requires that when we receive "the same input", we update on the same proposition E, and the E:not-E Bayes factor of our updates is the same.

In other words, once it's settled what the input is, commutativity leaves no further room for sensitivity to one's prior credence in underminers.

Now there are ways out of this too. Some of them are the same ways we mentioned before. We could give up Assumption-2. We could give up the idea that the world *settles* which input we receive without consulting our current credence in underminers. So even if you and I have the same experiential phenomenology, it might be appropriate for our credal models to be supplied different inputs. But the most controversial element in the preceding is (iv).

The third argument dispenses with (iv), and also with (iii). In their place it puts the following fact:

- (vi) Given a E:not-E Bayes factor, and an initial probability function $\text{Old}(\cdot)$, your ending credence in E will then be a function solely of $\text{Old}(E)$. *It will not be sensitive to the credence $\text{Old}(\cdot)$ assigns to anything else—in particular, not to $\text{Old}(U)$ —except insofar as that affects the factors already specified.*

Claims (i), (ii), (v) and (vi) tell us that if you and I have the same initial probability in E, and acquire the same further immediate justification to believe it, then we should *end* with the same probability in E, too. Even if I have

more justification to believe some underminer of the immediate justification we acquired than you do.⁶³

Again, this is not tolerable. Something here still needs to give. Getting rid of (iii) and (iv) didn't really help.

10. Conclusion

I will close by summing up some of the stable packages that do still look viable for a dogmatist or credulist. They're all built out of pieces we've already mentioned.

Option 1. We could limit the explanatory ambitions of our probabilistic frameworks even further than is commonly acknowledged, and say they just don't apply in cases where there is the prospect of non-quotidian undermining. Or at least, they don't apply fully. As we saw, we can try to model non-quotidian undermining, imperfectly, as if it were a case of quotidian undermining.

We might explore other formal systems to see if they can better model the behavior of non-quotidian underminers.

Option 2. We might stay with Bayesian systems, and say that our prior credence in non-quotidian underminers plays a role in determining what the *inputs* to these formal systems are. The fact that a range of subjects acquire the same prima facie justification, informally understood, doesn't guarantee that our formal models of their epistemic state should all be updated in the same way. This gives up Assumption-2 and -3.

Option 3. We mentioned the possibility of your *sometimes* updating on bold claims, like Hand. In Classical Bayesianism, we'd have to assume "Popperian" conditional probabilities, so that the credence might later be defeated. But reflecting on what happens if you *acquired the undermining evidence first* pushed us to include at least some elements of Option 2. In Jeffrey Bayesianism, we can immediately make sense of you updating on claims like Hand, but we'd have to reconcile ourselves with the impossibility of pure underminers, in the sense envisaged by Assumption-6. We'd also have to have to include elements of Option 2 here, too, because of difficulties raised in the second and third arguments from section 9. We'd need to have some response to the worries

⁶³As Christensen points out, this is arguably the same difficulty that Garber 1980 presses against Field. Garber phrases his complaint in terms of intuitive redundancies, rather than undermining effects, that Field's proposal is not respecting. But formally, it's natural to expect those effects to stand or fall together.

Wagner 2002's response to Garber is different from the recommendation (iii) described in note 60. I would describe the move he recommends here as: don't assume that what update you should perform is determined just by what experiential phenomenology you have, or what prima facie justification you newly acquire. It is also sensitive to facts about what you've previously learned. This amounts to giving up our Assumption-2 and Assumption-3. Hawthorne 2004a responds similarly to Garber.

raised in the second argument about how to measure the amount of support we acquired for what we update on.

Option 4. We might say we always update on claims that are *more cautious than can possibly be undermined*—if there are any such claims, which I myself doubt. This strategy can still make room for dogmatism or credulism if Assumption-2, and possibly some other assumptions, are rejected. Section 7 above explored some ways this might be developed.

Bibliography

Peter Achinstein, *The Book of Evidence* (Oxford University Press, 2001).

William Alston, *Epistemic Justification* (Cornell University Press, 1989).

Paul Boghossian, “Knowledge of Logic”, in C. Peacocke and P. Boghossian, eds., *New Essay on the A Priori* (Oxford University Press, 2000), 229–54.

Paul Boghossian, “How are Objective Epistemic Reasons Possible?” *Philosophical Studies* 106 (2001), 1–40.

Paul Boghossian, “Blind Reasoning,” *Aristotelian Society Supplementary Volume* 77 (2003), 225–48.

Laurence Bonjour, *In Defense of Pure Reason*, (Cambridge University Press, 1998).

Alex Byrne, “Introspection,” *Philosophical Topics* 33 (2005), 79–104.

Alex Byrne, “Knowing That I Am Thinking,” in A. Hatzimoysis, ed., *Self-Knowledge* (Oxford University Press, 2008), 105–24.

Roderick Chisholm, *Theory of Knowledge*, 3rd ed, (Prentice-Hall, 1989).

David Christensen, “Confirmational Holism and Bayesian Epistemology,” *Philosophy of Science* 59 (1992), 540–57.

David Christensen, “Epistemology of Disagreement: The Good News,” *Philosophical Review* 116 (2007), 187–217.

David Christensen, “Does Murphy’s Law Apply in Epistemology? Self-Doubt and Rational Ideals,” *Oxford Studies in Epistemology* 2 (2008), 3–31.

David Christensen, “Higher-Order Evidence,” *Philosophy and Phenomenological Research* 81 (2010a), 185–215.

David Christensen, “Rational Reflection,” *Philosophical Perspectives* 24 (2010b), 121–40.

David Christensen, “Disagreement, Question-Begging and Epistemic Self-Criticism,” *Philosophers’ Imprint* 11 (2011).

- Eli Chudnoff, "The nature of intuitive justification," *Philosophical Studies* 153 (2011), 313–33.
- Stewart Cohen, "Why Basic Knowledge is Easy Knowledge," *Philosophy and Phenomenological Research* 70 (2005), 417–30.
- Stewart Cohen, "Bootstrapping, defeasible reasoning, and a priori justification," *Philosophical Perspectives* 24 (2010), 141–59.
- Juan Comesaña, "On an Argument Against Immediate Justification," in M. Steup and J. Turri, eds., *Contemporary Debates in Epistemology*, 2nd ed. (Blackwell, forthcoming).
- Fred Dretske, "Conclusive Reasons," *Australasian Journal of Philosophy* 49 (1971), 1–22.
- Adam Elga, "Lucky To Be Rational," (ms available at <http://www.princeton.edu/~adame/papers/bellingham-lucky.pdf>).
- David Enoch, "Not Just a Truthometer: Taking Oneself Seriously (but not Too Seriously) in Cases of Peer Disagreement," *Mind* 119 (2010), 953–97.
- Richard Feldman, *Epistemology* (Prentice-Hall, 2003).
- Richard Feldman, "Respecting the Evidence," *Philosophical Perspectives* 19 (2005), 95–119.
- Richard Feldman, "Epistemological Puzzles about Disagreement," in Stephen Hetherington, ed., *Epistemology Futures* (Oxford University Press, 2006), 216–36.
- Richard Feldman, "Evidentialism, Higher-Order Evidence, and Disagreement," *Episteme* 6 (2009), 294–312.
- Hartry Field, "A Note on Jeffrey Conditionalization," *Philosophy of Science* 45 (1978), 361–67.
- Bas van Fraassen, *Laws and Symmetry* (Oxford University Press, 1989).
- Richard Fumerton, *Metaepistemology and Skepticism* (Rowman and Littlefield, 1995).
- Daniel Garber, "Field and Jeffrey Conditionalization," *Philosophy of Science* 47 (1980), 142–45.
- James Hawthorne, "Three Models of Sequential Belief Updating on Uncertain Evidence," *Journal of Philosophical Logic* 33 (2004a), 89–123.
- John Hawthorne, "Deeply contingent a priori knowledge," *Philosophy and Phenomenological Research* 65 (2002), 247–69.
- John Hawthorne, *Knowledge and Lotteries* (Oxford University Press, 2004b).
- Michael Huemer, "Direct Realism and the Brain-in-a-Vat Argument," *Philosophy and Phenomenological Research* 61 (2000), 397–413.

- Michael Huemer, *Skepticism and the Veil of Perception* (Rowman and Littlefield, 2001a).
- Michael Huemer, "The Problem of Defeasible Justification," *Erkenntnis* 54 (2001b), 375–97.
- Michael Huemer, "Phenomenal Conservatism and the Internalist Intuition," *American Philosophical Quarterly* 43 (2006), 147–58.
- Michael Huemer, "Compassionate Phenomenal Conservatism," *Philosophy and Phenomenological Research* 74 (2007), 30–55.
- Michael Huemer, "The Puzzle of Metacoherence," *Philosophy and Phenomenological Research* 82 (2011), 1–21.
- Richard Jeffrey, "Alias Smith and Jones: the Testimony of the Senses," *Erkenntnis* 26 (1987), 391–99.
- Thomas Kelly, "The Epistemic Significance of Disagreement," in T. Gendler and J. Hawthorne, eds., *Oxford Studies in Epistemology*, vol 1 (2005), 167–96.
- Thomas Kelly, "Peer Disagreement and Higher Order Evidence," in R. Feldman and T. Warfield, eds., *Disagreement* (Oxford University Press, 2010), 111–74.
- Thomas Kelly, "Disagreement and the Burdens of Judgment," in D. Christensen and J. Lackey, eds., *The Epistemology of Disagreement: New Essays* (Oxford University Press, forthcoming).
- Matt Kotzen, "Silins's Liberalism," *Philosophical Studies* 159 (2012), 61–68.
- Matt Kotzen, "A Formal Account of Epistemic Defeat," (ms available at http://matthewkotzen.net/matthewkotzen.net/Research_files/defeatersweb.pdf).
- Peter Kung, "On Having No Reason: Dogmatism and Bayesian Confirmation," *Synthese* 177 (2010), 1–17.
- Marc Lange, "Is Jeffrey Conditionalization Defective By Virtue of Being Non-Commutative? Remarks on the Sameness of Sensory Experience," *Synthese* 123 (2000), 393–403.
- Matthew McGrath, "Dogmatism, Underminers and Skepticism," *Philosophy and Phenomenological Research* (forthcoming).
- Ram Neta, "Liberalism and Conservatism in the Epistemology of Perceptual Belief," *Australasian Journal of Philosophy* 88 (2010), 685–705.
- Alvin Plantinga, *Warrant: The Current Debate* (Oxford University Press, 1993).
- James Pryor, "The Skeptic and the Dogmatist," *Nous* 34 (2000), 517–49.
- James Pryor, "Highlights of Recent Epistemology," *British Journal for the Philosophy of Science* 52 (2001), 95–124.
- James Pryor, "What's Wrong with Moore's Argument?" *Philosophical Issues* 14 (2004), 349–78.

- James Pryor, "When Warrant Transmits," in A. Coliva, ed., *Wittgenstein, Epistemology and Mind: Themes from the Philosophy of Crispin Wright* (Oxford University Press, forthcoming).
- James Pryor, "Hypothetical Oughts," (ms available at <http://www.jimpryor.net/research/papers/Hypothetical.pdf>).
- Geoff Pynn, "The Bayesian Explanation of Transmission Failure", *Synthese* (forthcoming).
- Hans Reichenbach, *The Direction of Time* (University of California Press, 1956).
- Wesley Salmon, *Scientific Explanation and the Causal Structure of the World* (Princeton University Press, 1984).
- Joshua Schechter, "Rational Self-Doubt and the Failure of Closure", *Philosophical Studies* (forthcoming).
- Stephen Schiffer, "Skepticism and the Vagaries of Justified Belief," *Philosophical Studies* 119 (2004), 161–84.
- Nicholas Silins, "Basic Justification and the Moorean Response to the Skeptic," in T. Gendler and J. Hawthorne, eds. *Oxford Studies in Epistemology*, vol 2 (2008), 108–42.
- Nicholas Silins, "Experience and Defeat," *Philosophy and Phenomenological Research* (forthcoming).
- Scott Sturgeon, "Pollock on Defeasible Reasons," *Philosophical Studies* (forthcoming).
- Chris Tucker, "Why Open-Minded People Should Endorse Dogmatism," *Philosophical Perspectives* 24 (2010), 529–545.
- Chris Tucker, "Movin' on Up: Higher-Level Requirements and Inferential Justification," *Philosophical Studies* 157 (2012), 323–40.
- Carl Wagner, "Probability Kinematics and Commutativity," *Philosophy of Science* 69 (2002), 266–78.
- Carl Wagner, "Is Conditioning Really Incompatible with Holism?" *Journal of Philosophical Logic* (forthcoming).
- Brian Weatherson, "Induction and Supposition," *The Reasoner* 6 (2012) 78–80; available at <http://www.kent.ac.uk/secl/philosophy/jw/TheReasoner/vol6/TheReasoner-6>
- Ralph Wedgwood, "A priori bootstrapping," in A. Casullo and J. Thurow, eds., *The A Priori In Philosophy* (Oxford University Press, forthcoming).
- Jonathan Weisberg, "Commutativity or Holism? A Dilemma for Conditionalizers," *British Journal for the Philosophy of Science* 60 (2009), 793–812.

Jonathan Weisberg, "Updating, Undermining, and Independence", (ms available at http://www.utm.utoronto.ca/~weisber3/new/Research_files/Updating%20Undermining%20and%20Independence.pdf).

Roger White, "Problems for Dogmatism," *Philosophical Studies* 131 (2006), 525–57.

Tim Willenken, "Moorean responses to skepticism: a defense," *Philosophical Studies* 154 (2011), 1–25.

Timothy Williamson, "Knowledge and scepticism," in F. Jackson and M. Smith, eds., *The Oxford Handbook of Contemporary Philosophy* (Oxford University Press, 2005), 681–700.

Timothy Williamson, *The Philosophy of Philosophy* (Blackwell, 2007).

Timothy Williamson, "Improbable Knowing," in T. Dougherty, ed., *Evidentialism and its Discontents* (Oxford University Press, 2011), 147–64.

Timothy Williamson, "Very Improbable Knowing," (ms available at http://www.philosophy.ox.ac.uk/_data/assets/pdf_file/0015/19302/veryimprobable.pdf).

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

Stephen Yablo, "Coulde, Woulda, Shoulda," in T. Gendler and J. Hawthorne, eds., *Conceivability and Possibility* (Oxford University Press, 2002), 441–92.