

# Merits of Incoherence

jim.pryor@nyu.edu

for Pittsburgh Conference on Perceptual Experience  
and Empirical Reason, Oct 2016

## Section 1

1. My main interest when doing epistemology is in the conditions, nature, and “logic” of a status or quality that folk language may have no unambiguous direct expression for. Amongst fellow theorists, I can direct your attention to this status by calling it “prospective justification or warrant to be more confident” that something is the case.

2. I say “justification or warrant” because I make no subtle distinctions between these, as some other authors do. Another way I’ve sometimes captured this dimension of the quality is to say that it’s *epistemically more appropriate* for you to be more confident. But with all of these expressions, we need to draw on conventional examples to separate them from their intuitively close kin.

I say “prospective justification”, meaning thereby what some authors (including past versions of myself) call “propositional justification.” I’ve come to find the latter phrase unfortunate, because it can suggest that what’s being justified is (in the first place) a proposition, rather than an attitude like believing or suspending, or some kind of attitude dynamic, like becoming more confident. But that ought to be a substantive question, not one we hard-code an answer to in our shared vocabulary. (Also, I am sympathetic to the “attitude” answer rather than to the “proposition” answer.) Another thing the phrase “propositional justification” can suggest is that where the justification comes from is some set of propositions, those which constitute your evidence. Again, that is a substantive commitment, that our shared vocabulary ought not to decide prematurely. (Again, this is a commitment I would myself resist.)

I say “to be more confident” meaning to be more confident in some hypothesis than you should/would be in the absence of that justification. Later in this discussion we’ll encounter an opportunity to distinguish between having *more justification* to be confident in P, and having *justification to be more confident*. For the moment, I want to treat these indifferently; but I haven’t found a vocabulary that sounds neutral between them.

You'll notice that I talk about *confidence* rather than all-or-nothing doxastic stances like assenting or believing. It is controversial which (if either) of these kinds of doxastic stances we should think of as more fundamental. I have leanings; but I think we can side-step this controversy. Whichever of these kinds of stance is more fundamental, I assume it is feasible and can be productive to engage in theoretical inquiry about when one should be more confident, without also thereby inquiring into when it's appropriate for one to close inquiry or "settle" into an all-or-nothing doxastic stance. (In earlier work, I routinely spoke about "justification to believe," but my inquiries were primarily directed at the normative qualities of un-settled states like confidence.)

For ease of reference, let me abbreviate "prospective justification or warrant to be more confident" as PJC.

3. It's natural to divide the conditions that contribute to your having PJC for some hypothesis Q into two exclusive categories. (I don't know whether they are exhaustive.) Some of these conditions are the kinds of things we'd be willing to call "your justification for" Q, and which could make up "your grounds" for having some doxastic stance towards Q. (We're interested in these conditions even when *you don't* take any stance towards Q, but merely ought to; our question here is whether the conditions are such that, and then so related to you that, *you could* respond to them qua grounds.) Perhaps there are two questions here: whether the conditions merely *could be* your grounds, and whether you could *reasonably* so rely on them. But let's pass over that subtlety for now. The other category are conditions that may play some role in *making* you justified, but which theorists are more prone to call *enablers* of your having that status, rather than counting as "your justification" or being the kind of thing you could sensibly *base* a doxastic stance on. By PJC I mean only to include things from the first category: potential grounds, rather than mere enablers.

Epistemologists use the term "defeater" in several ways. Some use it to include considerations that work like negative enablers. This usage was especially prevalent in the Gettier literature, but is not confined to those discussions. On the other hand, many contemporary epistemologists, including myself, use "defeater" to mean a kind of potential ground.

Defeaters come in several varieties. One straightforward kind of defeat of your PJC for Q, would be some PJC for not-Q, or any hypothesis that's clearly incompatible with Q. This can be called "opposing" or "overriding" defeat, though we should not construe the phrases "overriding" or "defeat" here as success terms. The kind of defeat we're considering comes in degrees and can itself be defeated.

A more complex and interesting variety of defeat is what theorists have taken to calling "undermining" or "undercutting" defeat. Again, this can come in degrees and itself be defeated, so the labels should not be understood as success terms. I've argued (in "Problems for Credulism," 2013) that this category shouldn't be thought of as excluding the more straightforward one — there can

be “mixed”/hybrid cases — and also that we shouldn’t take it to be obvious that the two categories exhaust the universe of defeaters. It is very tricky to define or explicate what makes something an undermining defeater. I will take it for granted that we all understand this category well enough to inquire into it, and can agree about some paradigm examples.

One controversial issue we’ll be considering below is whether higher-order PJC — that is, PJC for some (perhaps false) claim about the epistemic facts in general or about your own epistemic status — can be defeaters at all for the corresponding first-order issues. For example, if you have some good first-order grounds to be confident that Q, and then acquire PJC for the (false) claim that you don’t (and never did) have such grounds, does that on balance justify you in being less confident that Q? If higher-order PJC can have this kind of defeating effect, I will assume it counts as a kind of undermining defeat. This is a substantive assumption, though many others make it too. There may be important differences between this kind of undermining defeat, and more commonplace examples, such as evidence that your instruments/senses aren’t reliable about Q. I don’t mean to preclude such differences. I assume only that there are also important similarities, and that we can productively theorize about them to some extent as a single class.

(On the other hand, some philosophers like Wedgwood argue that all undermining defeaters work via higher-order mechanisms. I’ve sometimes been tempted to this view, but I want to leave it open here.)

Undermining PJC may have a positive counterpart (e.g., evidence that your instruments/senses are more reliable than you thought). We could call this *bolstering* PJC.

4. I’ve defended an epistemology of perception I called “dogmatist.” I wasn’t the first to defend the kind of epistemology in question, and nor was I the first to use that name in a philosophical discussion. Though in doing so, I started a thread of using the name for more-or-less the kind of epistemology I defended. My use of “dogmatist” was inspired by its original application to Stoic epistemologists, who thought, against their Skeptic contemporaries, that it was possible for us to have reasonable beliefs. There are of course other usages of “dogmatist” and its cognates in philosophy: there’s the pejorative use, and also a use to refer to the Kripke/Harman puzzle about whether evidence for Q entitles one to expect any future evidence against Q to be misleading. There’s also Richard Jeffrey’s use of “dogmatic” to refer to theories that prioritize all-or-nothing belief over degree-like doxastic attitudes. Let’s ignore all these other usages, and confine our attention to the thread of using “dogmatist” for more-or-less the kind of epistemology I defended. A lot of variation in usage still remains, which it might be useful to sort out.

Before I explain my understanding of “dogmatist,” though, and how it differs from some others’ usage, let’s consider a similar terminological thread. This concerns the pair of terms “liberal” and “conservative”. Again, these have

other usages in philosophy that are irrelevant to our concerns here. The term “conservative” even has another usage in epistemology, to mean a “negative coherence theory” that says attitudes count as being justified, or absolved of the need for justification, just by being had. The thread of using these terms in epistemology that I’ve participated in (I think I started this one, too) is different. As I understand these terms, they are exclusive but not exhaustive positions a theorist can hold about a chosen epistemic flaw or vulnerability. To illustrate, suppose Jessie has what is or at least seems to be PJC for the hypothesis H that she has hands. Now attend to considerations or hypothetical evidence such that, were Jessie to become so related to it that it became a potential ground for her, it would undermine the PJC she has for H. Call such considerations or evidence U. What epistemic relations does Jessie need to have towards U, as part of her having PJC for H? The liberal position says that Jessie can have PJC for H without needing any antecedent PJC against U. She doesn’t need any reason to “rule out” U in advance. Nor does U (or its threatening content) need to in fact be false. But if Jessie does go on to *learn* or acquire U as evidence, or get PJC for it, that will undermine/defeat the PJC she otherwise, and until then, had for H.

The conservative position, on the other hand, says that Jessie *does* need to have PJC against U, as part of having PJC for H in the first place. She may be able to rule out U because of some “default entitlement,” rather than from any inquiry she engaged in. But without PJC against the potential underminer U, Jessie can’t have PJC for H in the first place. Some forms of this view might add that U also has to in fact be false, for Jessie to have PJC for H. Perhaps she even has to *know* not-U, to have PJC for H. But I understand those as stronger forms of conservatism, not part of the label’s meaning.

These views are not exhaustive. A third position will say that U does have to *in fact be* false, for Jessie to have PJC for H, though she doesn’t need antecedent PJC to believe it it’s false. If U is “My senses are unreliable,” then reliabilists about perception provide a good example of this third view. They say that Jessie will have (perceptual) PJC for H only if U is false, that is, only if her senses are in fact reliable. She doesn’t need evidence that her senses are reliable. And even if her senses *are* reliable, evidence that *they aren’t* (that is, for U), will defeat the PJC she has for H.

There may be yet further positions. When undermining evidence comes to light, some epistemologists are inclined to say that it shows that the subject never really had justification for H in the first place. I’m not sure whether this is best understood in terms of the third position I described, or as something new.

I want to emphasize that as I understand these positions, they are *not monolithic*. One might be liberal about some epistemic vulnerabilities, conservative about others, and have the third position about yet others. I’m not sure that it’s coherent to be liberal about all vulnerabilities. In my view, one kind of possible underminer for my PJC for H is the claim that I lack such PJC. Evidence that I lack such PJC would defeat (contribute towards defeating) that PJC, even if I

in fact had it. In this case, surely the right attitude to take to the underminer is the third position. For me to have the PJC, *this* possible underminer really does need to be false. It's not enough for me to merely lack evidence to believe it's true.

I'm not sure that it's coherent to be conservative about all vulnerabilities, either. But this is a position that some epistemologists are attracted to.

5. Now let's return to the term "dogmatism." As I said, I advocated and defended an epistemology of perception that I applied that label to. At some point, though, I realized that the commitments that I was taking to be definitive of "dogmatism" were shared by a large family of views, where I had just been speaking in favor of one specific subspecies of the family. The specific view I advocated was distinguished by being internalist, concerning our PJC for observations about our perceptual environment, and giving a central explanatory role to the phenomenology of perceptual experience. Other views that differed in these respects could also be "dogmatists", as I came to understand that term. To give just two examples, one could also call oneself a dogmatist if one thought that PJC (or whatever epistemic status one was working with) required one's senses to in fact be reliable, or if one thought that PJC was only acquired when one genuinely perceived, not when one hallucinated.

Around the same time I was advocating my version of "dogmatism," Michael Huemer was defending a view of much the same sort, that he called "Phenomenal Conservatism." That label applies more specifically to the subspecies of dogmatism that he and I are most sympathetic to, rather than to the whole family. Huemer also defended this view more broadly than just for perception. I am sympathetic to some generalizations (for example, to memory and to math), but I am also open to there being other sources of justification, which he resists. We don't agree about all the arguments we each give for our views, but of course there is some overlap.

Of course, Huemer and I weren't the first to defend dogmatist views. In modern days, roots of the view can be seen in the common sense philosophy of Reid and Moore, and the particularism of Chisholm. (How much ground we share with Chisholm is an interesting issue, whose answer isn't obvious. See Alston 1997 for some relevant details.) More recently, you can see views more recognizably like ours in Pollock 1974 and in a series of "modest foundationalisms" that were developed in the 1970s and 1980s. These views were clearly members of the broad dogmatist family, as I understand it, though many of them weren't Huemer and my favored *kind* of dogmatism. Pollock for example gave a central role to a neo-Wittgensteinian theory of concepts, and others of our predecessors gave a central role to a non-propositional notion of acquaintance. Also, not all of our predecessors here were foundationalists *about perception*: some thought we had immediate but defeasible justification only about internal subject matters.

I do believe in immediate but defeasible justification, but I don't count myself as a foundationalist, because that requires additional commitments about what

other sorts of justification are or aren't possible, and about the global structure of justification.

Since the mid-90s, many other dogmatist views have been appearing. Sometimes the specific sympathies of these views are like Huemer's and mine, but others diverge in crucial respects (such as by not being internalist). I hope that Huemer and I contributed some to these developments, but the momentum behind them was certainly already out there in the field when we arrived.

6. In the field, not everybody uses the term "dogmatism," nor the terms "liberal" and "conservative," in the ways I've explained them here. Some use "dogmatist" to refer to the specific views Huemer and I advocated about perception. Some associate that term with specific commitments about the legitimacy of Moore-style proofs about the external world. I have myself argued that Moore-style proofs aren't guilty of all the vices they've been charged with (though they're no saints either). And I think there are natural motivational connections between dogmatism and the arguments I gave there. But I regard these connections as substantive, not part of the definition of "dogmatism." I'd call some views "dogmatist" that rejected my views about Moore-style proofs. But as I said, some other authors do understand "dogmatism" to include extra commitments (which I myself undertake) about Moorean arguments.

It's common also to see the term "liberal" used interchangeably with "dogmatist." What I would myself say is: being a dogmatist commits you to non-conservatism about every potential underminer. But one may take a liberal position wrt some underminers, and the third position wrt to others — in fact this may be the only coherent view, overall.

If you're reading any given author who uses the terms "liberal," or dogmatist," or some nearby term like "neo-Moorean," and you want to know specifically what they mean by it, you'll have to look closely at what that author specifically says and/or assumes. There seems to me to be too much variation in the details of different authors' usage for us to have a shared exact understanding of these labels.

7. In Pryor 2013, I proposed the label "credulism" for a family of views that was broader than dogmatism. (The motivating idea was that being credulous was somewhat like being dogmatic, though perhaps not so extreme.) In terms of the notions explained above, credulism can be identified with the view that we should be non-conservative about some potential underminers. A credulist is allowed to say we should be conservative about others. For example, one might plausibly think that a subject who has arithmetic justification for some result can have that justification be undermined by evidence that she's taken a bad-at-math-pill, without thinking any antecedent reason to think she *hadn't* taken such a pill needed to be part of her original justification. That would count as a credulist view of this species of justification. It permits one to affirm or to deny the further claim that any arithmetic justification is immediate. (Perhaps arithmetic justification always requires one to have antecedent PJC for, say, the

principle of non-contradiction.)

Whereas only a few epistemologists endorse (a) the specific views Huemer and I advocated about the epistemology of perception, a much broader group endorse (b) views that fit my broad understanding of “dogmatism” (including dogmatic accounts of justification other than perceptual), and even many more endorse (c) views that fit my understanding of “credulism.” Indeed, among theorists with explicit commitments, anti-credulists are in a small, polarized minority. Even BonJour when he was a coherentist was a credulist (and arguably even a dogmatist, about our justification for claims about what we believe).

Curiously, most of the objections I’ve encountered to (a), or at any rate, the ones that my opponents seemed to be most moved by, really seemed to most fundamentally be challenges for (c). My specific views (a) may well be wrong, but if they are, I think it’s not going to be for these reasons. I and most other epistemologists alike have to find some way to meet these challenges to credulism. This paper explores one such pressing challenge, and how we might respond to it. For concreteness, and in the theme of this conference, I will draw on dogmatist epistemology of perception as an example target. But I think that fundamentally this is really an issue for any credulist.

You may well find my response to this challenge a cure that’s worse than the disease.

8. We’ll get to the challenge, and my response to it, only in the closing Section 9 of this paper. In between, we need to introduce and coordinate our assumptions about a range of general issues about justification, undermining evidence, and incoherent attitudes.

## Section 2

Consider Matthew, who has just proved some theorem T.

Consider also Noma, who seems to herself in the same way to have proved a theorem T’. That is, Noma has all the superficial experiences of having proved T’. But in her case, let’s suppose this is an illusion. She’s made a subtle mistake. You are welcome to suppose that T’ is or isn’t really provable; I require only that Noma isn’t aware of a proof of it, but only seems to be. Perhaps she even *has produced* a proof, but her understanding of it is critically flawed in ways she’s not aware of. So it merely feels to her like she’s successfully grasping a proof of T’.

In some such cases Noma may have inductive grounds for thinking that T’ probably has been proved. Perhaps whenever she seems to have proved something, it tends to be true. Or depending on her track record, maybe she’ll have inductive grounds for believing the opposite! For the present discussion, though, I want to set any such inductive grounds aside. Let’s suppose Noma lacks any inductive

evidence about the track-record of her mathematical seemings. She just has the experience itself, of seeming to have proved and understood the proof of  $T'$ .

I'll want to refer to some versions of these stories where that is all the evidence Matthew and Noma have. Call the subjects in those scenarios Matthew-0 and Noma-0.

I'll also want to refer to versions of these stories where Matthew and Noma get *additional*, "higher-order" evidence, telling them that they aren't justified in their conclusions. For example, perhaps their brilliant mathematical uncle looks over their work and tells them it isn't a proof. It's important for these versions of the stories that the new evidence really does justify Matthew and Noma in doubting that they had a proof, and doubting they were justified in believing their respective "theorems." In Matthew's case, though, those claims are false. He did have a proof, and before his uncle showed up, at any rate, he *was* justified in believing his theorem. We can call Matthew's new evidence *misleading* because it's evidence for false claims. Nonetheless it does give him justification for those claims. Call the subject in that scenario Matthew-M. ("-M" for "misleading (higher-order, defeating) evidence.")

In Noma's case, the uncle as I've described him may be *telling her something true*. Noma may in fact not have a proof. But we can imagine versions of the story where Noma's uncle is still *misleading* her, because he's just making up a criticism, and doesn't really believe she lacks a proof; and also versions where her uncle is in fact pointing out to her exactly what the mistake in her initial reasoning was. Let's call these subjects Noma-M and Noma-Revealed.

I'll also want to refer to versions of these stories where Matthew and Noma don't have any higher-order *evidence*, but they do have negative higher-order *attitudes* — attitudes they (arguably) ought to have formed in response to the evidence described before. They may have these negative higher-order attitudes because of a crippling insecurity. On some epistemological views, they may have them in response to *illusions* of acquiring the -M evidence. But it's important to these versions that Matthew and Noma merely believe — *while lacking justification for believing* — that they don't have a proof and aren't justified in accepting their "theorems." In other versions they may merely doubt this, or be agnostic about it. Call all such subjects Matthew-A and Noma-A. ("-A" for "akratic" or "agnostic.")

It will make discussion easiest if we assume that in the -M cases, Matthew and Noma *do also form* the higher-order attitudes described in the -A cases. (Some of our discussion should carry over to the cases where they don't.) It will also make discussion easiest if we assume that in *all* the cases, Matthew continues to believe  $T$  on the basis of his (genuine) proof, and Noma continues to believe  $T'$  on the basis of her (illusory, or ill-understood) proof. In some of the versions, then, they are believing in ways that their higher-order critical beliefs tell them they shouldn't. (Some of our discussion should carry over to cases where Matthew and Noma don't do this.)



Here is a table of these possibilities:

	genuine proof	illusion of proof
no additional evidence	Matthew-0 (uninteresting)	Noma-0
gets negative higher-order evidence	Matthew-M	Noma-M  Noma-Revealed (uninteresting)
has negative higher-order attitudes	Matthew-A	Noma-A

I label Matthew-0 “uninteresting” because it seems straightforward what our epistemological theories should say about him. He has and understands a proof; so presumably he’s justified in believing on that basis the theorem he’s thereby proved. Of course there will be different explanations of the nature of that justification; but what I’ve said so far won’t be contested.

I label Noma-Revealed “uninteresting” because it also seems straightforward what our epistemological theories should say about her. She seemed to grasp a proof, but an informant she has reason to trust has explained to her what her mistake was; so presumably she’s not justified in continuing to believe her “theorem.” Again, there will be different explanations of how and why that works. But what I’ve said so far won’t be contested.

The remaining cases: the -M and -A cases, and Noma-0, are more interesting and provoke a variety of theoretical positions.

We might loosely arrange those positions on a spectrum, with “more objective” views on one end, and “more subjective” views on the other end. (It shouldn’t be supposed, though, that all accounts of these cases vary along only one dimension, and can be linearly ordered between those endpoints.) Very subjective views will tend to say:

S1. Noma-0 is just as justified in believing her “theorem” as Matthew-0 is in believing his.

S2. All the agents with higher-order attitudes (whether based on evidence or not) that say they are in a deluded, Noma-like condition, are unjustified in believing their “theorems.”

Very objective views will tend to say:

O1. Only Matthew (-0, -M, and -A) has any justification for his “theorem”; Noma-0 has none. (Except perhaps for inductive justification, which we’re here ignoring.)

O2. What's important for Matthew's justification to believe T is *just* whether he does in fact grasp a proof. So he's entitled to *ignore* the higher-order evidence or attitudes that mislead by saying he doesn't. At least, he can ignore them when it comes to the question of T's truth.

I'm not happy with, or willing to straightforwardly accept, any of these four claims. I think the truth is somewhere in the middle of this loose spectrum, and will require some care to articulate (much less to convincingly establish).

Rather than "objective," sometimes that end of the spectrum is instead called "externalist" (and the other end "internalist"). This can be confusing, since there's no direct, definitional connection between the kind of "objectivity" described here and externalism in the sense associated with Goldman, nor the sense associated with Williamson.

Sometimes all the views not at the "objective" (or "externalist") end of the spectrum are called "subjective" (or "internalist"). Other times all the views not at the "subjective" end are called "objective". These terminological practices are also confusing.

I have another, specific reason for discomfort with the term "subjective," as some hear this to mean that *it's the subject's perspective* that determines her normative relation towards a proposition Q. But on many views that oppose the most "objective" end of the spectrum, the subject's higher-order evidence and attitudes needn't determine her normative relation towards Q all by themselves. They may *contribute to* making the subject stand in some such relation, but they won't trump other contributing factors, and arguably there can be other contributors that are unrecognized by the "objectivist" than a subject's higher-order evidence and attitudes.

For all these reasons, I dislike and prefer to avoid the qualifiers "objective" and "subjective" (or "external" and "internal"). These terms do more harm than good in these discussions. But it is helpful to have in mind the loose arrangement of views I've used those terms to introduce.

With respect to S1 versus O1: I'm inclined to deny that Noma-0 has *nothing* by way of justification. But I'm also inclined to doubt that she's *epistemically on equal footing* with Matthew-0. There seems to be *some* positive epistemic relation she stands in to T', but *something better* about the relation that Matthew stands in to T. (Even if T' and T do both happen to be theorems.) I wish I understood this more, and were in a position to say more about it (or even knew what were the right questions to ask, to get us moving forward). But I don't. Nonetheless, this is how I'm inclined to regard these subjects, and in discussing these cases with others, I've found that many others, perhaps most others, are also so inclined. But there is no consensus.

The rest of our discussion will focus on what to say about the -M and -A cases, for both Matthew and Noma.

## Section 3

The past twenty years in ethics and epistemology have seen a flourishing of work around a few questions:

Q1. What are the normative effects of mere (unjustified) attitudes? This work appears under labels like “wide-scope oughts” and “structural normativity,” and one of its central contributors is John Broome, but there are also many other labels and contributors.

Q2. What are the normative effects of higher-order information? In epistemology, one of the central contributors to this work is David Christensen, but there are also many others.

These first two questions interact with each other, when we consider cases like Matthew-A and Noma-A who have mere (unjustified) higher-order beliefs, rather than higher-order evidence. Also, much of the past decade’s work on “the epistemology of disagreement” has engaged with aspects of both Q1 and Q2.

Continuing our questions:

Q3. Can it ever be reasonable (or justified or rational) to be inconsistent, akratic, or so on?

Later, I’ll introduce an general umbrella term “incoherent,” and will restate Q3 as: Can it ever be reasonable to be “incoherent”?

Finally:

Q4. Are “normative tragedies” or “dilemmas” or “rational toxicity” possible? that is, cases where the subject has no option she can reasonably take?

After developing my combination of answers to these questions, I’ll talk through what it tells us to think about our -M and -A cases.

Ask a dozen different ethicists or epistemologists their views on the four issues listed, and nowadays you’re likely to hear a dozen (or more) different combined accounts. I cannot hope to survey all the candidates, or to do more than begin to motivate the combination I find most workable. But as we walk through these questions, I’ll gesture towards some other popular accounts besides the ones I’m recommending.

## Section 4

Let’s begin with questions about the normative effects of mere (unjustified) beliefs. Can such beliefs *justify us* in believing their obvious inferential consequences? One extreme view (from the “objective” end of Section 2’s spectrum) will say these beliefs have no normative effect at all. Only what justification (or: what

justified beliefs) you have matters. The other, “subjective” end of the spectrum will say that justification is just a matter of the downward dynamics from your subjective perspective, that is, the beliefs you start with. So “mere” beliefs aren’t in themselves handicapped as justifiers.

Much of the discussion of these issues has taken place within a debate between those who advocate “narrow scope” and those who advocate “wide scope” construals of intuitive principles like:

If you believe P, and P obviously entails Q, then you ought to believe Q (or: then believe Q!)

and:

If you believe you ought to believe Q, then you ought to believe Q (or: then believe Q!)

This debate also engages with several principles of *practical* reason, but I’ll confine our focus to epistemology here.

We’ll return to principles like the second in Section 5, below.

Observe that in the antecedents of these principles, mental states are specified that we’re not assuming to be justified. The “narrow” interpretation of the principles sees them as having the form:  $A \rightarrow \text{Ought } B$ , and so *merely having* the A states suffices for it being true that you ought to form, or are justified in forming, the belief specified in the consequent. The narrow-scooper’s opponents complain that this predicts more “Oughts” or justification than really exist in such cases. For instance, if I have a stupid belief that P, that by itself shouldn’t make me justified in believing  $P \vee P$ . Neither is  $P \vee P$  something I ought, all things considered, believe. I *shouldn’t* believe it, *nor* the premise P which entails it.

The “wide” interpretation of the principles sees them as having the form  $\text{Ought}(A \rightarrow B)$ , and rejects the inference pattern of “factual detachment” that would permit us to infer, from the additional premise that A is merely true, that Ought B. Some who take this path are sympathetic to the inference pattern of “deontic detachment,” which says that from  $\text{Ought}(A \rightarrow B)$  and the additional premise that *Ought* A (that is, you’re *justified* in having the A states), Ought B does follow. But this is contentious, so let’s leave it open.<sup>1</sup>

A wide-scooper can occupy a position somewhere in the middle of Section 2’s spectrum. As we said, they will represent our intuitive principles as  $\text{Ought}(A \rightarrow B)$ . They will represent a mere unjustified attitude as a case where A but not Ought A. They can give the combination of A and  $\text{Ought}(A \rightarrow B)$  some

---

<sup>1</sup>How you evaluate these interpretations will be affected by whether you think of the original principles as merely saying something like: a *necessary condition for* being reasonable (or justified or rational) is that if ... then ... Or whether you think of the original principles as aiming, more interestingly, to *explain or identify the source of* what makes ... reasonable for you. Let’s follow most authors, who understand these principles in the second way.

interesting normative role; but they won't give it the significance that Ought B. For example, if you have an unjustified belief in P which obviously entails Q, the wide-scooper denies that you're thereby *justified* in believing Q, but can allow that you do stand in some kind of interesting normative relation to Q, of a sort to be unpacked.

A wide-scooper doesn't *have to* say this. They can instead occupy a position more towards the "objective" end of the spectrum, and say that in such cases there aren't *any* interesting normative relations present. In practice though, those who have embraced the most "objective" end of the spectrum have tended to reject the wide-scooper's proposals, and have *also* rejected the narrow-scope principles set out above. They'd instead only accept narrow-scope interpretations of *other* principles, such as:

If you're *justified in* believing P, and P obviously entails Q, then you ought to believe Q.

They'll say these principles apply to you so long as you *in fact are* justified in believing P, regardless of whether you know or believe that you are.<sup>2</sup>

I'm sympathetic to the view, advocated by the first group of wide-scoopers, that if some states A would, *when you had justification for them*, ground or justify other, inferentially "downstream" attitudes B, then even when you lack justification,

---

<sup>2</sup>One assumption shared between those who identify as "wide-scoopers" and those who identify as "narrow-scoopers" is that "Ought" in the above principles can scope differently wrt a conditional sentence like "If you believe P (and . . .), then believe Q." This is a substantive assumption, that we ought not to accept uncritically.

Many linguists nowadays favor a different view, associated most prominently with Angelika Kratzer, that says that "if"-clauses aren't antecedents of independent conditional sentences, but instead are restrictors for binary modal quantifiers. Consider the sentence:

Most students who study physics believe Q.  
= Most (student who studies physics) (believe Q)

On the Kratzer view, the role played by "Most" in that sentence can also be played by "Must" or "Ought", and when it is, clauses like "if you study physics" play the same role as "students who study physics":

If you study physics, you must/ought to believe Q.

really has a form like:

Must/Ought (studies physics) (believes Q)

The Kratzerians would analyze:

If you believe P, you ought to believe Q.

in the same sort of way. On these views, it's a mistake to think we that we have an independent sentence "If you believe P, then you believe Q" that offers multiple scopes (surrounding the consequent, or surrounding the whole conditional) for "Ought" to occupy.

A different challenge comes from Mark Schroeder, who argues for different reasons that the relevant, "deliberative" uses of "Ought" don't in fact take sentential complements. (His positive picture of how these sentences work is also different from Kratzer's.)

Some parts of the narrow-versus-wide debate can be reconstructed if these alternative views about the syntax of "Ought" sentences is accepted, but the *terms* in which we'd have to express the reconstructed views will look different. I will continue to talk about "scope" to stay connected with the existing literature, with which some readers may be familiar, but I doubt that this is really the best way to frame the fundamental issues.

having A still puts you in *some* kind of interesting normative relation to B. But not the relation of simply justifying B.<sup>3</sup>

Among those who recognize an interestingly different kind of normative relation here, we see a plethora of terminology, none of it fully satisfactory. The old, familiar kind of normative relation gets called “reasons rationality” or “evidential norms” (though that label doesn’t generalize well to the practical case, with which these discussions also engage). I’ve called it “categorical normativity,” with the idea of opposing it to “conditional” or “hypothetical normativity” on the other side.<sup>4</sup> The newer kind of normative relation posited here gets called “structural” or “attitudinal rationality,” or “normative” or “rational requirements.”<sup>5</sup> Some authors use terms like “coherence requirements.” There is a broad tendency to use terms like “rationality” and “coherence” in reference to this second kind of normativity, though other authors use “rationality” to refer to the whole genus of which we’re now distinguishing two varieties.

With reservations, I’ll use my labels of “categorical” versus “hypothetical” normativity. When you’re justified in some beliefs A that obviously support a conclusion B, I’ll say you have categorical justification for both A and B. When the justification for A is removed but you retain those (now unjustified) beliefs, then I’ll say that they merely hypothetically justify or support having belief B.

What I’m mostly interested in isn’t the relation of hypothetical support but rather the relation of hypothetical defeat, especially hypothetical undermining defeat. If your evidence E categorically supports believing Q, but justification for some undermining hypothesis U, if you had it, would undermine that justification for Q, then I say that *a mere belief* in U hypothetically undermines the categorical support E gives you for believing Q. That’s different than merely being a *possible* underminer: we’re saying *more* than just that U is capable of undermining, in the familiar sense, when you acquire justification for it. In order for U to hypothetically undermine, you need to have some definite doxastic attitude towards U, such as believing it. Suspending judgment in U may also have some undermining effect, though presumably not as severe as outright belief in U would have. Having a doubt that U may also have some negative effects, but these are even weaker (and more elusive).

An especially interesting variety of hypothetical defeat comes from *higher-order* hypothetical underminers. If higher-order *evidence* that your grounds E don’t support believing Q can (categorically) undermine the justification that E actually does give you for Q, then the mere (unjustified) higher-order *belief* that E doesn’t support believing Q will hypothetically undermine belief in Q. Mere (unjustified)

---

<sup>3</sup>Our intuitive principle, above, talks only of the case where the content of A *entails* the content of B; but the phenomenon seems to be more general.

<sup>4</sup>I didn’t mean to prejudge issues about whether it’s categorical in the Kantian or Footian senses, that is, applied to agents regardless of their contingent properties.

<sup>5</sup>Broome uses the last label. He initially used “normative requirements” but then revised his terminology so as not to prejudge the question of whether it was genuinely normative, that is, whether we *had any reason to comply* with these requirements. Broome’s use of “requirements” in his labels still does prejudge some questions about *the structure* of these normative relations.

*suspended judgment or doubts* about whether E supports believing Q will have similar, but weaker, effects.

Calling some attitudes “hypothetical” supporters or underminers is just to stick a label on them, and to affirm that they exhibit some interesting normative property. What’s the cash value?

Here are some claims held in common by different accounts of these “hypothetical” normative relations:

- (1) They characterize a kind of goodness/badness that holds between **combinations of mere** (not-necessarily categorically justified) attitudes.
- (2) We reject a principle of **factual detachment** for this goodness. From the fact that beliefs A hypothetically support belief B, and that you have beliefs A, it does not follow that you’re categorically justified in having B. In other words, from Ought(A  $\rightarrow$  B) and A, it does not follow that Ought B. (The status of “deontic detachment,” which concerns what happens when Ought A, is as I said contested.)
- (3) This is a **distinctive kind** of goodness/badness. Suppose Huey has an unjustified belief, and isn’t yet sure whether to accept some obvious consequence of it. The original belief and its consequence are both disconfirmed by his actual evidence. Dewey has the same evidence and unjustified belief, but he accepts the obvious consequence. Louie has the same evidence and unjustified belief, and he rejects the consequence. Dewey is the most coherent of these subjects, and seems thereby to be exhibiting a virtue the others aren’t. Further, Huey and Louie’s dissent<sup>6</sup> from the consequence seems to exhibit some further, distinctive failing, additional to their failing in having the original unjustified belief. (In Louie’s case, he exhibits this failing in part *because* he conforms to his evidence in rejecting the conclusion.)

Other claims about hypothetical normativity are more disputed.

- One question is how hypothetical and categorical relations compare. Which of Huey, Dewey, and Louie is doing best, epistemically? Dewey is the most coherent. But Louie’s belief conform most closely to his evidence, and Dewey’s the least. I won’t propose any straightforward ranking of these subjects, though other proposals I’ll make do bear on this.<sup>7</sup>
- Another question is how hypothetical and categorical relations agglomerate. My view is that hypothetical and categorical norms can combine to generate more hypothetical norms. That is, if attitudes A hypothetically support

---

<sup>6</sup>I’ll say you *dissent* from a conclusion when you either disbelieve it or deliberately withhold belief from it.

<sup>7</sup>In the epistemology literature, the dominant tendency is to emphasize categorical normativity, and thus if forced to rank these subjects, to rank Louie first. In the practical literature, there’s a widespread (but not universal) tendency to instead to emphasize coherence, and thus to rank Dewey first. Worsnip (PQ, 2016?) discusses why the fields might diverge in this way.

believing B, and your evidence is such as to make B categorically support C, then — even if you don't in fact infer from B to C — that can be enough to make A also hypothetically support C. But your actual attitudes towards B and C may interfere with this in complex ways.

- Some wide-scopers discuss cases where there's a body of attitudes A1..An that each seem to be justified, but you seem unjustified in holding them jointly. For example, if we're permissivists about some variety of justification, perhaps choice B and choice not-B might be equally OK, but choosing both not be OK. Quinn's self-torturer gives another example of this. I agree that cases of hypothetical badness can also exhibit this form, where some attitudes only exhibit badness when held jointly. But unlike some theorists, I'm not convinced that everything exhibiting this form is a case of hypothetical badness. There may be facts about what combinations of attitudes one is categorically justified in having that aren't a function of the categorical justification of each attitude taken individually. But I can't defend this stance here.
- When it comes to questions about what kind of "prospective" or "propositional" categorical justification you have, I think the hypothetical normative facts are irrelevant. Prospective justification is insulated from the merely hypothetical.
- But I do think hypothetical facts matter when it comes to whether your belief is well-founded or "doxastically" categorically justified. If your evidence E categorically supports believing Q, but that support is hypothetically defeated by some (perhaps unjustified) attitudes D, then so long as you retain those attitudes D, you can't justifiably believe Q on the basis of E. (If you see yourself to have hands, but believe your senses are unreliable, then even if that belief is unjustified, you can't then have a reasonable perceptual belief that you have hands.)<sup>8</sup>

This last point bears on a question asked by Tim Willenken. If a subject has prospective justification to believe Q, must there be available to her some epistemically permissible way to form that belief? Willenken thought the answer had to be yes. I think the answer may be no. It can happen that there's some doxastic response that your epistemic position makes appropriate, but that there's no epistemically good dynamic route from your current attitudes to that response. (Buridan's Ass may be a helpful analogue: the ass has reason

---

<sup>8</sup>This view about how hypothetical justification matters to well-founded belief means that the simple picture of well-founded belief in B as:

- (i) you have prospective/propositional justification for B,
- (ii) you do believe B, and
- (iii) your belief is based on the justification described in (i)

is inadequate. I'm not envisaging that your belief in U has to make it impossible for your belief in B to be based on your evidence E, which does in fact prospectively justify it. So we need a more complex story about the relation between well-founded belief and prospective justification. (I'd hope that a more satisfying story could be told than just adding a fourth condition to the simple picture, but I don't know what that more satisfying story is.)



to be walking to the left or to the right, but getting himself to move in one of the directions rather than the other has to be non-rational.) Why? Because you may also have (unjustified) beliefs in U that get in the way of having a doxastically justified belief in Q. And it might not be permissible for you to *refrain* from believing Q, either, since after all believing Q is what your evidence does support.

Wouldn't it in such cases always be permissible for you to *give up* your unjustified belief in U, and upon doing so, *then* form a justified belief in Q? I don't know. Just because your belief in U is unjustified, it's not obvious that you'd have any other justified attitudes upon which you could *base* your change of mind about U. If you're lucky, you'd *come to realize* that U is not supported by your evidence, and then you may base a withdrawal from U upon that realization. But I'm trying to theorize about what's the reasonable doxastic responses for imperfect epistemic agents like us, who haven't always recognized all the facts about their epistemic situations, and thus need to make doxastic choices even in advance of having such realizations. If that insight about U hasn't yet occurred to you, it's not obvious whether any change of your mind about U could be well-formed. Of course, neither is your belief in U well-formed. If you had never believed U in the first place, but suspended judgment about it all along, that attitude could presumably have been well-formed.

I wish I understood better how well-formedness works for suspended judgment, and what is the relation between well-founded changes of mind, and the availability of well-founded attitudes at the end of a change in mind. In advance of having that all worked out, I'm sympathetic to the idea that in many cases like the one just described, you'd be in a dynamic dilemma: one where no change of mind could be well-founded, but some dynamically *un*justified changes of mind may *end with* you having well-founded attitudes afterwards. We'll discuss this further in Section 6, below.

## Section 5

Now let's turn to questions about how first-order evidence interacts with higher-order evidence (and attitudes).

One question is: can higher-order evidence (regarding how much justification you have to believe Q) have downward effects, that is, contribute to or make you have less (or more) justification for the first-order belief Q? The stock example here is that you've done some first-order, mathematical reasoning that seems to (and let's suppose, really does) support Q, but then you get evidence that you've been drugged in a way that makes you bad at math, in ways that it's hard for you while drugged to otherwise detect.

Williamson's Clock is a case where you can know certain imprecise facts about the position of a pointer on an unmarked clock face. We suppose you can't

know the precise position, because your belief about that wouldn't be sufficiently "safe" to count as knowledge (you'd still have that belief were the pointer in an ever-so-slightly different position). But you can sometimes know that the pointer *is contained in region r*, since when the pointer is in the middle of that region, the nearby worlds where the pointer's position differs but you still have the belief that the pointer is in r would be ones where it's elsewhere in the same region, and so that belief is still true. However, can you *know that* you know the pointer is in r? If we suppose that r is the smallest region where it's possible for you to have the first-order perceptual knowledge that the pointer is in r, then that knowledge would only be possessed in the case where the pointer is exactly in the middle of r. So to *know that* you know the pointer is in r, arguably you'd have to know that the pointer is exactly in the middle of r. (At least, you'd be brought to this conclusion if you reasoned in the way we're doing here, and accepted the premise that you did know the pointer was in r.) Since you can't know precise facts like that about the pointer's position, it seems you *can't* know that you know the pointer is in r, though the result still stands that you can have the mere first-order knowledge, that the pointer *is* in r. This shows that  $K(R)$  is consistent with  $\neg KK(R)$ . Williamson goes on to argue that  $K(R)$  is also consistent with your knowing that it's *very unlikely* (indeed, more or less *arbitrarily unlikely*) that you  $K(R)$ . So  $K(R)$  doesn't preclude having even a large amount of  $J\neg K(R)$ .

Williamson's discussion focuses on knowledge, but similar reflections might tempt one to say that  $J(R)$  is also consistent with your knowing (or being justified in believing) that it's unlikely that you're  $J(R)$  (see Christensen, PP 2010). If correct, that would show that  $J(R)$  doesn't preclude your having  $J\neg J(R)$  either. Other strong bridge principles between higher-order and first-order justification might be similarly challenged. The only principle I'm sympathetic to in this neighborhood is the weak claim that  $J\neg J(R)$  constitutes *some negative justification towards R*. It tends to *disconfirm R to some extent*. Gaining that higher-order justification contributes towards your being *less* justified in believing R than you'd be if you lacked it.

We might also have sympathy for the idea that higher-order justification can have *positive* downward effects. For example,  $JJ(Q)$  may give you some positive justification towards Q. David Barnett has an interesting case that challenges this idea, or at least demands that the idea be refined: You have a body of first-order evidence about a murder, and are not certain about what conclusions it supports. Some oracle (Sherlock Holmes, or another source of trustworthy information about your epistemic status) surveys your evidence and tells you, "Actually, some of your evidence has the property of decisively supporting conclusion Q but as a matter of fact I know independently that the evidence in question is false." That is, the oracle is telling you you have some false but justified belief (the belief is not specifically identified) that confirms the conclusion Q. The oracle does not tell you whether Q is in fact true, or whether Q is supported by the oracle's larger body of evidence, nor even what conclusions you should draw after receiving the information he gave you. Many people David has told

this case to (myself included) judge that the agent is *not* then more justified in believing Q than she was before talking to the oracle. Yet it appears she did acquire further justification, that she earlier lacked, to believe she was justified in believing Q. So it seems that JJ(Q) does not always straightforwardly suffice for getting more justification for Q.<sup>9</sup>

## Section 6

Next let's take up the questions of "incoherence"; and also of "normative tragedies" or "dilemmas."

Sometimes holding inconsistent attitudes seems like it may be the reasonable thing for a subject to do. Consider (Cases-1) Preface Scenarios, where you've acquired good evidence for each of many claims, but you also have general evidence that you're likely to have made a mistake in your assessment of so many different questions. Can't it then be reasonable to believe each of the many claims, and reasonable also to simultaneously disbelieve their conjunction?

Depending on your circumstances, it might in some cases not be reasonable to *disbelieve* the conjunction, but merely to *withhold* belief in it. If you believe each of the many claims and merely dissent (but don't necessarily disbelieve) their conjunction, it's not clear we should call that "an inconsistent set of attitudes"; but I'll be interested in the question of when this can be reasonable, too.

Or consider (Cases-2) Long Deductions, where you have good evidence for an initial believed premise, from which you deduce a long sequence of lemmas, arriving at some conclusion that really is entailed by the initial premise. However, acknowledging the accumulating likelihood of error across such a long chain of deductions, and aware of your own fallibility, you don't yet accept the entailed conclusion. Can you be reasonable in doing so?

In response to (Cases-1), many philosophers will advise us to move to a notion of *graded confidence*, rather than all-or-nothing belief. Perhaps that move can also help with (Cases-2). But as many authors have observed, analogous cases can arise for graded confidence, too. Suppose (Cases-3) I have good evidence that supports having credence 0.8 in P. I'm also confident that P is either logically equivalent to Q or logically equivalent to not-Q; but I'm not sure which. So I'm not sure whether my credence in Q should be 0.8 or 0.2. Aware of my own fallibility, I split the difference and have credence 0.5 in Q. In that case, I'd have credences that are probabilistically incoherent. But can't I be reasonable in doing

---

<sup>9</sup>We've been discussing the question whether higher-order justification can have downward impacts (whether positive or negative), and thus affect how much justification you have for your first-order beliefs. A different question is whether the first-order justification can have upward impacts. Many participants in the disagreement literature effectively say "no," that would license objectionably question-begging demotions of your peers (when you happen to be the one who gets things right). But Tom Kelly argues that the answer should be "yes," and my sympathies are aligned with his.

so? Or perhaps I do something more sophisticated: instead of simply having a credence of 0.5 in Q, I “split my opinion” into two probability distributions, in one of which I give Q credence 0.8 and in the other of which I give it credence 0.2. In that case, my total opinion would be represented by a set of probability distributions, some of which are probabilistically incoherent. But can’t I be reasonable in having that opinion?

Changing gears somewhat, consider (Cases-4) where I have good evidence E for believing Q, but then acquire some additional “higher-order” evidence H that justifies me in believing (falsely) that E does not support Q. What would I be justified in believing in response to my new body of evidence, E+H? Could it ever be reasonable in such cases for me to believe Q, while simultaneously doubting or agnostic whether I am justified (or whether E justifies me) in so believing?

There’s a definite feeling of awkwardness about the combination of attitudes described in these cases. One wants to say, “That can’t be a good position to be in, where you’d hold such an awkward body of opinions.” Even so, though, perhaps it can be the best way of doxastically responding to the unfortunate epistemic position one *is* in, in the cases described. That at any rate is the possibility I’ll be exploring. I’ll use the term “incoherent” as a general umbrella term to cover all the kinds of awkwardness illustrated in these cases: deductive inconsistency, a recognized lack of deductive closure, probabilistic incoherence, various forms of epistemic akrasia (such as believing Q while believing you lack justification to do so, or refraining from believing Q while believing your justification obliges you to believe Q). I don’t know that there is any natural genus of which these forms of awkwardness are all species. But they do in my mind have a family resemblance, and enough of the issues and other discussions we’ll be engaging with interact with several of them at once, that it’s useful to have a single term for them. Unlike some other authors, my use of the term “incoherent” is not meant by definition to imply unreasonableness, irrationality, or anything like that. Our substantive question is precisely whether it can ever be reasonable (or justified or rational) to be incoherent in any or all of these forms.

Instead of answering that question directly, I want to more cautiously propose that sometimes having attitudes that are incoherent in some of these ways is the “least bad” (or “most permissible”) doxastic response to one’s body of evidence and other attitudes. For example, the least bad response to Preface cases may be to have inconsistent beliefs. The least bad response to knowing that some claim is either tautologous or contradictory, but not knowing which, may be to have an intermediate credence, though then your credences would be incoherent. And so on. I say “least bad” to leave it an open, unsettled question whether we should say that (in at least some such cases) the incoherent set of attitudes can be outright *justified*: whether they can be *positively reasonable* ones for you to have. Or, on the other hand, whether they (in at least some such cases) exhibit a normative dilemma/tragedy: that is, a case where any doxastic choice you

make is bound to be objectionable or justificatorily defective, and so no response is outright justified. (This is not the same as *skepticism*; the skeptic will say that the response of suspending judgment is outright justified.)

One reaction to this proposal might be: “You shouldn’t be satisfied with the incoherent set of attitudes, you should reflect harder and figure out which attitudes are least supported by your evidence and change them.” Perhaps you should! But what review processes you should add to your mental to-do list is one question; and what doxastic responses you’re justified in forming here-and-now, on the basis of the insights you’ve already had (or should have had) is a different question; and I’m only intending to theorize about the latter.

Another reaction might be: “In such cases, why don’t you just *suspend* all the beliefs, then you won’t any longer have incoherent attitudes?” Well, it’s not clear on some probabilistic models what such suspending amounts to, and in particular not clear whether suspending would preserve coherence. But even bracketing those concerns, it’s just wishful thinking to think that suspension will always be the way out of these difficulties. Consider this case. Mathematicians have good but inconclusive reasons to believe that  $P < NP$ .<sup>10</sup> Now suppose you allegedly prove that some of the arguments mathematicians have for believing  $P < NP$  are flawed, apparently supporting the all-things-considered response that we should suspend judgment whether  $P < NP$ . But now you get evidence that you were under the influence of a “bad at math” drug when proving what you did. In this case, perhaps you should suspend judgment about *some* questions (such as whether your proof was correct), but it seems that the question at issue — whether  $P < NP$  — is one you’re *not* anymore justified in suspending judgment about.

What this case illustrates is that suspending judgment is also a doxastic response that one can adopt for certain grounds as opposed to others, and this response is vulnerable to undermining just as belief and disbelief are. So the fundamental issues we’re discussing bear on the rationality of suspending judgment too. There’s no guarantee in cases of conflict of the sort we’re considering that suspending judgment will always be the best, or a unproblematic, option.

I’ve invoked the idea of a normative dilemma/tragedy. A practical example of this may be if I’ve promised you to commit some evil act, and now the moment of truth is at hand. I shouldn’t have made that promise, but arguably, having made it, I am now normatively criticizable if I break the promise and refrain from committing the act. On the other hand, I would also be criticizable (we can suppose, more criticizable) if I follow through. So whatever I do here will be bad in some way. It’s controversial whether such cases can arise in the practical domain, and even more controversial whether they have any epistemic analogues.

---

<sup>10</sup>This means that the problems of complexity class P (solvable by a deterministic machine in time proportional to a polynomial of the input size) is a proper subset of the problems of complexity class NP (solvable by a *non*-deterministic machine in polynomial time, or alternatively, such that *solutions can be verified* by a deterministic machine in polynomial time).

Dilemmas/tragedies don't mean that anything goes: some options may be decisively off the table. Neither do they mean that all of the unexcluded candidates are equally bad. Some option(s) may be "least bad"; it's just that all options, including the "least bad" ones, still include *some* "unexcused" normative defect. (I don't like using the word "unexcused" here, since the issues isn't primarily about blameworthiness. But I don't know what word to use in its place.)

If some option is least bad, then isn't it guaranteed to be justified? Isn't it always reasonable to take the best option that's available? I don't know. That's a substantive normative claim. Many are sympathetic to it, but it's intelligibly deniable. If you did accept that claim, then yeah, you wouldn't think that dilemmas/tragedies of the sort I'm envisaging are possible.

I don't think that *incomparability* by itself generates a dilemma/tragedy of the sort I'm envisaging. It may just entail that the justificatory facts are indeterminate. (Or perhaps that justificatory statuses aren't linearly orderable.)

Arguably, dilemmas may include cases where your total evidence not only jointly recommends some attitudes that can't *coherently* be simultaneously adopted, but also stances that can't *in fact* be simultaneously adopted. (For example: could your evidence ever recommend {Believe P!, Believe Q!, Don't both believe P and believe Q!})

In discussing these issues, I've found it helpful to identify two thoughts that tend to go unquestioned (and usually aren't even explicitly articulated) in epistemology, but that I think are in fact quite substantive and intelligibly disputable. Thought 1 is about the Guaranteed Existence of some appropriate doxastic response. (There *is* an explicit debate about the Guaranteed Uniqueness of such a response.) That is, for any body of evidence, will there always be *at least one* doxastic response which that evidence justifies? (The answer to this question may depend on what menu of doxastic responses you countenance.)

A thought related to this (call it Thought 1\*) is that if some given doxastic response (such as suspending judgment) has its justification defeated, then some other response(s) must get their justification strengthened. If Thought 1 sometimes fails, presumably Thought 1\* can also fail.

A different idea is Thought 2: that if some attitudes are "incoherent" relative to each other, then your evidence can't justify the joint response of holding all those attitudes simultaneously. Depending on what view we take about reasonable incoherence, and/or about dilemmas/tragedies, we may want to reject or revise this principle.

I am only putting the thoughts on the table, and saying I don't want us to accept them uncritically. Some of the possibilities I'm exploring push against them. But I am not going to argue for a specific resolution here.

## Section 7

Consider Dede-0, who comes to be justified in the normal way in believing the consequence Q of some other justified belief P she has. And Dede-M, who in fact undergoes the same reasoning or insights that Dede-0 does, but also has evidence that ought to undermine that transition. Yet Dede-M ignores the undermining evidence and goes ahead and believes Q anyway, on the same grounds that Dede-0 does. I invite you to share my intuition here that Dede-M has done something wrong, and more specifically, that her final belief in Q is not epistemically appropriate and so is ill-formed.

If there are Closure Principles for well-founded belief, they need to be formulated in a way that tracks the presence or absence of undermining evidence, of the sort that Dede-M should be respecting but doesn't.

One currently popular way to formulate this principle is in the Williamson/Hawthorne style:

If you JBp and competently deduce Q from that belief in P, while retaining your justification for P, then the belief in Q so formed will be well-founded.

(Williamson and Hawthorne defend such principles for Knowledge, not justified belief.)

Perhaps in such a principle the notion of “competently deduce” can incorporate such things as: not ignoring evidence that undermines the deduction, as Dede-M did. I suspect that that is not the right place to include reference to such undermining evidence; it should instead be a separate clause. (I offer no argument for this suspicion.) More interestingly, when it's genuine *undermining evidence* at issue — and not just *mere beliefs* that undermining possibilities obtain — then I think such evidence should be as destructive of prospective/“propositional” justification as it is of well-founded belief. So even Closure Principles about prospective justification should include some reference to the absence of such undermining evidence. (In these cases, there may be fewer undermining possibilities to worry about, since the subject needn't have performed any deduction whose execution could be impugned. But there may still be scope for some undermining evidence, e.g., a philosophical argument that entailment in *that* logical system is irrelevant to epistemology.) One way to accommodate such undermining evidence is for the Closure Principle to merely have the consequence that you have prima facie justification to believe the entailed proposition, where that justification is yet vulnerable to being undermined.

## Section 8

So how does this all apply to our mathematicians from Section 2?

I assumed that Noma-0 has *some* justification for believing what she's seemed to prove, albeit not justification of the same kind and quality as Matthew-0 has. (I don't want to say, nor to deny, that Matthew-0 is justified in *being more confident* in his conclusion than Noma-0 is in hers.)

When it comes to Matthew-M and Noma-M, who have genuine (though misleading) undermining *evidence*, I've proposed that the subjects are less justified in believing their "theorems." This is because their higher-order justification contributes towards undermining those beliefs.

When it comes to Matthew-A and Noma-A, who merely have negative higher-order *attitudes*, I've proposed that their beliefs in their "theorems" are *hypothetically* undermined, and so cannot be well-founded or doxastically justified. But they can still be prospectively/propositionally justified.

Though Matthew-A and Noma-A may not be in a position to have well-founded beliefs in their conclusions, they may not be in a position to have well-founded dissent in those conclusions either. We discussed this kind of "dynamic dilemma" at the end of Section 4.

## Section 9

Our attitudes towards possible underminers raise a challenge for proponents of immediate justification for Q.<sup>11</sup> The puzzle starts from the idea that suspended judgment about an undermining possibility is something of an underminer, too, or at least some kind of obstacle to all-things-considered justification, even if it's less of an obstacle than *outright belief in* the underminer. So consider a case where you allegedly have immediate justification to believe Q, that would be undermined by evidence that U. In some such cases, an agent may *avoid having any* attitude (justified or not) towards U: perhaps because they never considered U, or they aren't conceptually sophisticated enough to consider it. But that won't be so *in general*. In many cases, the agent will at least *have justification* to have some attitude towards U. (We count deliberately suspending judgment about U as an attitude.) And the problem is that it seems like any justified attitude towards U other than disbelief is going to undermine the immediate justification, at least to some degree. Likewise, any mere (unjustified) attitude other than disbelief is going to hypothetically undermine. So it appears that the only way that the immediate justification for Q can survive without being undermined to at least some degree is when it's accompanied by disbelief in (or justification to disbelieve) U. That is the challenge. The fact that the justification is immediate might hold out the promise that you can get away without needing justified disbelief in U, but if the justification for Q is to survive, you can't.

---

<sup>11</sup>Really this challenge could be pressed more broadly against any Credulist, but I will just discuss this simpler version.



This is an interesting set of issues, in part because I think the correct response is different for the hypothetical side than for the categorical. The hypothetical issues are more complex, but dialectically it will work better to start with them. Here's one way to put the hypothetical challenge. Suppose you do believe Q on the basis of grounds that (you know) are vulnerable to being undermined by U. Then there's a *coherence constraint* on you to disbelieve U if you have any attitude towards it. You couldn't be justified in believing Q in that way and, say, suspending judgment about U. Compare this passage from Crispin Wright:

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

What I want to say in reply is: I agree that believing Q on those grounds imposes a coherence constraint on you to disbelieve U if you have any attitude towards it. But that is not the same thing as an *epistemic dependence* on disbelief in U. All it tells us is that a subject who believes Q in that way while having some other attitude towards U is doing *something* wrong. It doesn't follow that anything is wrong *with their belief in Q*. (And nor, if our earlier discussion of dilemmas/tragedies is right, does it follow that that isn't the best doxastic option for the subject. More on this in a moment.)

In more detail: Suppose the subject does disbelieve U, but without having any evidence for doing so. Then her belief about U will be unjustified; but is it obvious that her belief in Q on the envisaged grounds would also be faulty? No, this is not obvious. If the belief in Q *did* epistemically depend on disbelief in U, then the unjustified disbelief in U wouldn't discharge that debt. Unjustified beliefs can't justify other beliefs that epistemically depend on them. But if the belief in Q *doesn't* so depend, then the subject who has an unjustified disbelief in U would only be doing something wrong wrt U, and not with respect to her belief in Q. Alternatively, if the subject took some other attitude towards U, let's consider what effect that mere attitude would have (postponing consideration of the subject's justification for it for the time being). On my view, the mere attitude of suspending judgment towards or believing U would hypothetically undermine her belief in Q to at least some degree. That would prevent her belief in Q from being well-founded. But it would not affect her prospective/propositional justification for Q. If we want the subject to have prospective/propositional justification for Q, it doesn't matter what actual attitude she takes to U, but only what justification she has towards U (which question we will take up below). If we want the subject to have a well-founded belief in Q, then she must either have no attitude towards U (which may be a precarious option) or she must disbelieve U. But we haven't yet seen a compelling reason to think that disbelief in U needs to be justified.

Finally, if we're willing to countenance cases of justified incoherence, or cases of epistemic dilemmas/tragedies, the terms of this discussion will have to be different. Believing Q on the basis of grounds vulnerable to U, while suspending

judgment about U, may be *incoherent*; but it no longer follows that it's not justified, or at least the best doxastic option.

What about the categorical questions? Isn't it still the case that the only way that immediate justification can survive without being undermined is when it's accompanied by *justification to disbelieve U*? Here I want to distinguish between two sorts of cases in which your epistemic position might fail to require disbelief in U but also fail to require belief. In one sort of case, you have lots of evidence bearing on the question whether U, and the balanced verdict of that evidence is mixed. One paradigm of this is when you know that U's truth depends on a objectively chancy process, and U has a chance of 50% of being true. In those cases, we'd all agree that your evidence justifies suspending judgment about U, and we can also agree that this undermines your justification for Q (to some middling degree; more justification to believe U would undermine more). The other kind of case is where you don't have any, or much, evidence about U. These are cases that Keynes would describe as "uncertain" rather than in terms of "risk." In these cases, perhaps the right thing to say is that your evidence doesn't support *any* attitude towards U, not even the attitude of suspending judgment. Or perhaps the right thing to say is that your evidence supports a range of attitudes, that includes suspending judgment but includes more positive and more negative attitudes as well. Or perhaps the right thing to say is that your evidence *does* support suspending judgment, but that this is nonetheless an interesting different species of having justification to suspend judgment. However one wants to do that theoretical classification, I think the best thing for a proponent of immediate justification (or more generally, a Credulist) to say is that having this kind of evidential relation to U *does not* undermine your justification for Q, to any significant degree. (Of course we want *also* to be able to say that a situation in which you have justification to *disbelieve* U is better.)

This was a complex response, and some parts of it may seem unpalatable to you. I believe most views are going to have to say something somewhat uncomfortable here. The impression that one can sail through these issues cleanly tends to be fueled by the thought that *suspending judgment* will always be a rational fallback, and/or by the thought that for certain notable instances of U, we might have *default epistemic entitlements* to disbelieve them. But, in response to the first thought, we've already stressed (with the P < NP case) that the issues we're discussing affect the rationality of suspending judgment, too. And in response to the second thought, sure maybe for some choices of U we can get justification to disbelieve them without doing anything to earn it. But can't *that* justification to disbelieve those Us be undermined by *other* Us?<sup>12</sup> I think that *eventually*, we're going to find some justification and some U that threatens it where you're not in any antecedent position to justifiably rule that U out. The only views that can promise to avoid these prospects seem to be an extreme form of Coherentism, or a traditional foundationalist picture where some of our justification is just

---

<sup>12</sup>For example, by a philosophical argument that there's no such thing as unearned justification? Though some philosophers have conjectured that unearned justification can't be undermined but only defeated.

not underminable. I and most other contemporary epistemologists find these options unattractive.

There's an interesting contrast between what I've recommended in the Closure discussion and in the present context. Earlier, I was imagining some undermining evidence or belief in U, and urging that we *take it more seriously as a threat* to the subject's justified belief in a deduced belief Q. In the present context, I'm saying instead that *the subject's existing belief in Q* hypothetically commits her to disbelieving U if she has any attitude towards it. I think there's some truth in each of these ideas, and real-life examples will inevitably involve some mixture of them. It's just for presentational reasons that I've focused only on one effect at a time.