

From Impossibility to Evidentialism?

Alex Worsnip

Forthcoming in *Episteme*

Penultimate version; please cite published version where possible

It's often said that it is impossible to respond to non-evidential considerations in belief-formation, at least in certain ways. While it certainly isn't impossible for non-evidential considerations to *influence* our beliefs, it seems that it is typically not possible for us to believe *directly and consciously on the basis of* non-evidential considerations. To take a typical example, if I offer you \$1,000,000 to believe that the number of stars in the sky is even, it seems that you would not be able to so believe on the basis of the fact that it's financially advantageous for you to do so. At most you might be able to engage in some indirect "plotting" to try to get yourself to form the requisite belief.

While it's controversial exactly what the relevant "impossibility" comes to – merely psychological, conceptual, metaphysical – so far we haven't said anything *normative*. But many philosophers think that the impossibility of responding to non-evidential considerations, however it is to be understood, provides grounds for accepting some further, normative thesis. Typically, they think it supports one or both of the following:

Evidentialism-Reasons. Only evidential considerations can constitute normative reasons for belief.

Evidentialism-Ought. One (all-things-considered) ought to believe p only if one has adequate evidence for p .¹

Some philosophers also make the further claim that *Evidentialism-Ought*² is a *constitutive* norm of belief (though, as we'll see below, it's not always fully clear what this claim amounts to).³ This might also be thought to be supported by the impossibility of responding to non-evidential considerations in belief-formation.

There are a variety of ways in which one might try to support these normative theses by appeal to the impossibility-claim. In this paper, I want to put pressure on these various attempts by raising a simple, yet so far as I know so far overlooked, problem for them. In brief, the problem is that it isn't true that one cannot (directly and consciously) respond, in belief-formation to considerations that

This paper concerns topics I've been thinking about for a long time, and draws on two now-abandoned papers. The first was presented at a conference at the University of Pittsburgh in 2012, where I benefitted from Adam Marushak's commentary. The second was presented at the University of California San Diego, UNC Chapel Hill, the University of Pittsburgh, and Yale, all in 2015. Finally, the present paper was presented at the 2019 Episteme Conference in Skukuza, South Africa, where I benefitted from Bob Beddor's commentary. I thank the audiences at all of these places for many helpful questions.

¹ I won't be concerned with versions of evidentialism that are stipulatively restricted to the *epistemic* ought, or *epistemic* justification (cf., e.g., Conee & Feldman 1985), since these are compatible with allowing that practical considerations provide reasons for belief, or bear on what you all-things-considered ought to, despite the difficulty of responding to them.

² Since *Evidentialism-Reasons* isn't strictly speaking a norm, it *a fortiori* can't strictly speaking be a constitutive norm.

³ Cf., e.g., Adler (2002b); Shah (2006: 494); Nolfi (2018).

don't *actually* constitute (good) evidence for the proposition under consideration; what is true, at most, is that one cannot (direct and consciously) respond, in belief-formation to considerations that one *oneself takes* to be evidentially irrelevant to that proposition. While this point is fairly obvious once stated, and is occasionally noted as a clarification in passing,⁴ its significance hasn't been appreciated, or so I'll argue. I'll suggest that once we take full account of this point, the arguments from the impossibility-claim to the evidentialist theses stated above, or to constitutivism about them, don't go through – at least not in the form that they have been developed in so far. That's not to say that the evidentialist theses are false; they may yet be true. But it is to suggest that some of the most popular and seemingly powerful arguments for them fail.

The plan for the paper is as follows. §1 is a grumpy survey-cum-rant about the general failure of many participants in the debate to attend to the distinction between what one's evidence *actually* supports and what one *takes* one's evidence to support, when formulating the impossibility-claim and related claims. It spells out the need to weaken the impossibility-claim accordingly. §§2-4 reconstruct three arguments from the impossibility-claim to evidentialist normative conclusions – respectively, from a motivational constraint on reasons, from ought implies can, and by inference to the best explanation – and suggests that all of them lose much of their power once we take account of the weakening of the impossibility-claim that I am urging. In the section concerning inference to the best explanation, a kind of alternative explanation of our inability of respond to non-evidential considerations emerges, appealing not to evidentialism but to a distinct *coherence constraint* on belief. §5 considers some worries one might have about that alternative explanation. §6 wraps up.

1. What the evidence *actually* supports vs. what one *takes* one's evidence to support: the need for weakening the impossibility-claim

Epistemologists shouldn't need reminding that, at least *pave* certain very radically subjective Bayesians, there is a distinction between what one's evidence actually supports and what one takes it to support.⁵ I might take the fact that I drew a tarot card showing The Hanged Man to be strong evidence that some terrible fate will befall me tomorrow, but in fact it isn't strong evidence of this. Moreover, not all failures to correct identify what one's evidence supports involve stupidity or superstition or gross, manifest irrationality in the way that this first example (on most precisifications) does. For example, *whatever* doxastic attitudes the evidence justifies having about various complex philosophical questions – like whether we have free will, or whether the mental is reducible to the physical – it is highly non-obvious that it does so, and some intelligent, thoughtful people are misidentifying what the evidence does support.⁶ Such failures also need not involve self-deception or lack of awareness of one's own mental states.

⁴ E.g., Adler (2002b: 29); Shah (2006: 497).

⁵ I'll stay neutral here on what 'taking' amounts to: whether it requires full-fledged belief, or something less. For the purposes of discussion of inference and basing, some have thought that taking need not be full-fledged belief (see e.g. Boghossian 2014: 6-11). On the other hand, the thought that it's structurally irrational to fail to believe what one *takes* one's evidence to support may rely on a construal of 'taking' and belief.

⁶ It might be that the evidence is "permissive" with respect to such difficult questions. But even if that were so, there are some smart people who think that evidence permits only one response, and so there are still people misidentifying what the evidence supports.

Given the obviousness – and clear epistemological importance – of the difference between what one’s evidence supports and what one takes it to support, it is odd that in debates about believing at will, the possibility of responding to non-evidential reasons, and the like, it seems to go out the window. So, for example, Jonathan Adler (2002a: 8) makes the astounding claim that “believing in opposition to one’s evidence is motivationally unintelligible”, when presumably he meant instead to make the less astounding (though still bold) claim that believing in opposition to *one’s own judgment* about the evidence is motivationally unintelligible.⁷ (His paper is about the (alleged) impossibility of epistemic akrasia, which is usually understood as involving defiance of one’s own *judgment* about what the evidence, not just a failure to believe what the evidence *in fact* supports.) Examples of people failing to believe what their evidence actually supports are clearly intelligible and prevalent. This phenomenon need not come about because of their willfully defying what they recognize the evidence to support (as in cases of epistemic akrasia), but can happen in a more anodyne way when people are just mistaken about what their evidence supports.

Other examples abound. Here are just a few:

- Nishi Shah (2003: 462) writes that “one particular belief-forming process, reasoning, is regulated *solely* by evidential considerations” (his italics). This isn’t right, at least if an “evidential consideration” is a consideration that is *actually* evidence rather than one that the agent *takes* to be evidence.⁸ Considerations that one takes to be, but aren’t actually, evidence for *p* can certainly influence one’s reasoning about whether *p*. At most it’s considerations that one doesn’t *take* to be evidence that can’t feature in reasoning.
- Paul Noordhof (2004: 247) writes that “if we are consciously attending to the question of whether *p*, we cannot help but make a judgment in line with what the evidence gives us grounds for believing true”. This isn’t right. Lots of people consciously attend to the question of whether *p* and end up with a judgment that is not in line with what the evidence gives them grounds for believing true; uncontroversially, this can happen when they miscalculate their

⁷ The blurb on the back of Adler’s (2002b) book – which I would expect, but don’t know, that Adler himself wrote – even more flatly claims that “belief in defiance of one’s evidence (or evidentialism) is unintelligible”. It presents this as a major conclusion of the book, but what the book could hope to be showing, at most, is again only that belief in defiance of what one *takes* one’s evidence to support is unintelligible.

Elsewhere (2002b: 29, 31, 34-5), somewhat more carefully, Adler clarifies that such claims hold only on “the condition of full awareness”. But he is rather unclear about what this condition requires full awareness *of*. On p. 29, he implies that it’s enough for full awareness that one is fully aware of one’s *assessment* of one’s reasons or evidence. But in that case, it isn’t true that it’s impossible to believe in defiance of one’s evidence even under conditions of full awareness, since one might be aware of what one’s assessment of the evidence *is*, without being aware that it is mistaken. Similarly, on pp. 34-35, he implies that the barriers to full awareness are things like “nonconscious influences and distraction” or “mental disturbance” (cf. also p. 73). But again, one can simply believe against one’s evidence because one has misassessed what one’s evidence supports; this need not involve nonconscious influences, or distraction, or mental disturbance.

⁸ Perhaps one might suspect that ‘evidential consideration’ can take the latter reading. But at another point in his article, Shah writes “my deliberation won’t count as belief-formation, and the conclusion I draw won’t count as belief, unless the deliberation is solely influenced by evidence for and against *p*”; this quotation doesn’t take the latter reading. Admittedly, Shah is by this point in the article in the middle of summarizing a view that isn’t his own view. But (a) he seems not to think that the quoted claim is *false*, but only that it can’t explain something he thinks needs explaining and is thus *incomplete*; and (b) even as a matter of stating a plausible version of the view he’s critically discussing, Shah ought to have said “solely influenced by considerations I *take* to be evidence for and against *p*”.

evidence. At most it's true that such people can't help but make a judgment in line with *their own evaluation* of what the evidence gives them grounds for believing true.

- Tamar Gendler (2008: 565) writes that “belief aims to ‘track truth’ in the sense that belief is subject to immediate revision in the face of changes in our all-things-considered evidence.” This isn't right. A change in my all-things-considered evidence may have absolutely no effect on my belief if I don't recognize it as evidence. At most, belief is subject to immediate revision in the face of *perceived* changes in our all-things-considered evidence.

The list goes on. You may think I'm being pedantic or uncharitable here, and that the quoted authors obviously *meant* to make the calmer, more careful claims about what one *takes* the evidence to support. But at the very least they are failing to be clear about the distinction I'm insisting on (and I find it genuinely puzzling why the same mistake keeps repeating itself). Moreover, as we'll see over the next few sections, brushing over the distinction has very real effects on the debate about evidentialism.

Before that, I want to note a slightly distinct, though related, debate where inattention to the distinction also causes substantive errors. In the debate about whether we have *responsibility* or *control* over our beliefs, it's often said that our beliefs are “controlled by the evidence,” “determined by the evidence,” or “at the mercy of the evidence”; belief is something that “happens to us, not something we do”; it's not “up to us.”⁹ That's supposed to set a puzzle for how we can have control over them or responsibility for our beliefs (and, to some, definitively shows that we don't have such control or responsibility. But the initial premise is wrong: our beliefs aren't just controlled or determined by the evidence, because the effect of our evidence on our beliefs is mediated by our evidential standards: what we take to be evidence for what. Plausibly, since our evidential standards belong to us, this grounds a sense in which our beliefs are (partially) products of our agency and are not just imposed from without by the evidence.¹⁰

The lesson of this section is just the one that I previewed in the introduction. We are going to be considering arguments for evidentialism that rely on the premise that it is impossible to respond (directly and consciously) to non-evidential consideration in belief formation. Taken on its most natural reading, ‘non-evidential considerations’ with respect to some proposition *p* are considerations that don't in fact constitute evidence for or against *p*. But read this way, the arguments will not be sound, because it is possible to respond directly and consciously to considerations that don't in fact constitute evidence for or against *p*, where one *takes* those considerations to be evidence for or against *p*. The only hope for these arguments is to employ a weaker impossibility-claim, which says (roughly) that it's impossible to respond (directly and consciously) to considerations that one *takes* to be evidentially irrelevant. The question to be investigated in the next three sections is whether weakening the impossibility-claim in this way allows the arguments to go through.

2. Arguments from a motivational constraint on normative reasons

⁹ See e.g., Williams (1973: 147-8); Owens (2000: 12); Flowerree (2017: 2768), Wedgwood (2017: 71).

¹⁰ See also Owens (2000: 25).

Thomas Kelly (2002) and Nishi Shah (2006) offer similar arguments for *Evidentialism-Reasons*. The argument goes roughly as follows:

Impossibility Claim. It is impossible to believe on the basis of non-evidential considerations.

Constraint on Reasons. Some consideration C constitutes a normative reason to Φ only if it's possible to Φ on the basis of C.

Therefore,

Evidentialism-Reasons. Non-evidential considerations cannot constitute normative reasons for belief (i.e., only evidential considerations can constitute normative reasons for belief).

In this statement of the argument, I have (following Kelly) formulated the argument in terms of what one believes *on the basis of*.¹¹ One might instead use the language of *motivating reasons*, or the reasons *for which* one believes.¹² The impossibility-claim would then be that non-evidential considerations cannot serve as motivating reasons for belief – the reasons *for which* one believes – and the constraint on (normative) reasons would be that some consideration is a normative reason for Φ -ing only if it can serve as a motivating reason for Φ -ing. I take this difference to be largely cosmetic, since I take the notion of a motivating reason for Φ -ing (or a reason for which one Φ 's) to be roughly equivalent to the notion of a consideration that one Φ 's on the basis of. In any case, any subtle differences between these notions shouldn't, I think, be very important for what follows.

Though the Kelly-Shah argument has been much-discussed,¹³ I think the most glaring weakness in the argument has been missed, which is that at least *prima facie*, the argument equivocates on 'epistemic consideration'. As we saw in the last section, *Impossibility-Claim* is only at all plausible if we take 'non-evidential considerations' to mean something like 'considerations that the agent *takes* to evidentially irrelevant'.¹⁴ Let's make that explicit by precisifying *Impossibility-Claim* a bit:

Impossibility Claim-Precise. It is impossible for A to believe p on the basis of considerations that A takes to be evidentially irrelevant to whether p .¹⁵

¹¹ In this I follow Kelly (2002).

¹² See the formulations of the argument given by Shah (2006) and Leary (2017). Shah's version of the argument effectively adds an ancillary argument for the impossibility-claim: he holds that C constitutes the reason for which one Φ 's only if it's possible to treat C as a premise in deliberation about whether to Φ . Since he thinks it is impossible to treat non-evidential considerations as premises in deliberation about whether to Φ , this yields the result that it's impossible for non-evidential considerations to be the reasons for which we believe.

¹³ See e.g. Steglich-Petersen (2008), Reisner (2009), Rinard (2015) and Leary (2017) for prominent responses to it. For my own part, I don't think that any of these responses succeed in undermining the argument, but it would take me too far off-track here to explain why.

¹⁴ Shah (2006: 497) more or less recognizes this. Having noted that it's possible to base one's belief upon something (including, say, a fact about one's own good) that one mistakenly *takes* to be evidence, he insists that this "doesn't falsify the claim that doxastic deliberation only weighs evidential considerations", which means he must be using 'evidential consideration' to mean a consideration that the agent *takes* to be evidence. But he doesn't recognize the ways, that I'm about to make explicit, in which this doesn't leave the argument where it was.

¹⁵ I don't mean to affirm that even this weaker claim is true. But I'll grant it, at least for now, for the sake of argument.

But in *Evidentialism-Reasons*, which is supposed to be the conclusion of the argument, the phrase ‘non-evidential considerations’ is most naturally interpreted as referring to considerations that aren’t *actually* evidence for p . After all, the evidentialist stance is that only (genuine) evidence for p constitutes a reason to believe p , not that only considerations *that the agent takes* to be evidence for p constitute a reason to believe p . Again, let’s make this precise:

Evidentialism-Reasons-Precise. Considerations that do not actually constitute evidence for p cannot constitute normative reasons to believe p (i.e., only considerations that do actually constitute evidence for p can constitute normative reasons to believe p).

So, once both of those claims have been made precise, we’re left with the following argument:

Impossibility Claim-Precise. It is impossible for A to believe p on the basis of considerations that A takes to be evidentially irrelevant to whether p .

Constraint on Reasons. Some consideration C constitutes a normative reason to Φ only if it’s possible to Φ on the basis of C .

Therefore,

Evidentialism-Reasons-Precise. Considerations that do not actually constitute evidence for p cannot constitute normative reasons to believe p (i.e., only considerations that do actually constitute evidence for p can constitute normative reasons to believe p).

This argument is patently invalid.

At first glance, what *Impossibility Claim-Precise* and *Constraint on Reasons* actually seem to entail is:

Subjective Evidentialism-Reasons. Considerations that A takes to be evidentially irrelevant to whether p cannot constitute normative reasons for A to believe p (i.e., only considerations that A takes to be evidence for p can constitute evidence normative reasons for A to believe p).

Could the evidentialist be contend with having established *Subjective Evidentialism-Reasons*? I think not; it is both too weak and too strong for the evidentialist’s purposes. Let me explain.

First, it is too weak for the evidentialist, because it doesn’t establish that some non-evidential consideration can’t be a normative reason for belief as long as the agent is *taking* it to be a reason. So for example, suppose that you’re offered \$1,000,000 to believe p , and suppose that you (falsely and irrationally) think that the offer of money is evidence that p – perhaps because you wrongly (falsely and irrationally) think that there’s an angel presiding over you, who ensures that you’re only ever offered money to believe true propositions. It seems that the argument does not rule out the fact that you’ve been offered money from counting as a normative reason for you to believe p .

The evidentialist might respond by pointing out that *Constraint on Reasons* on its own need not rule out all the non-reasons. In order for something to be a normative reason for you to believe p , it has to (i) be something you can believe p on the basis of (i.e., meet the *Constraint on Reasons*) and (ii) be

something that counts in favor of believing p . But for this reply to have any force, the evidentialist has to be presupposing that the fact that you're offered \$1,000,000 to believe p doesn't count in favor of believing p (or doesn't count in favor of it in the right way to be a reason to believe p). And that assumption is clearly question-begging in this context. The Kelly-Shah argument was supposed to provide a non-question-begging argument for evidentialism by arguing from a constraint on reasons that has independent motivation and plausibility. If that constraint isn't what's really doing the work in ruling out the offer of money from counting as a reason, then we've given up on the ambitions of the argument.

Second, and just as importantly, *Subjective Evidentialism-Reasons* is too strong (both for the evidentialist, and on its own intuitive merits). For it seems to rule out some consideration that *is* evidence for p , but which the agent takes not to be evidence for p , from constituting a normative reason to believe p . But then agents will be let off the hook for failing to respond to their evidence whenever they fail to recognize what it supports. For example, suppose that the Orioles' terrible losing record last year plus their lack of any free-agent acquisitions over the offseason together provide decisive support for believing that they will perform badly this year. But suppose also that, though I'm aware of the losing record and the lack of free-agent acquisitions, I don't recognize that the evidential force of these considerations, claiming that it's a new season and last year's results are irrelevant. If considerations that I don't *take* to be evidentially relevant to whether p are ruled out from being normative reasons for me to believe p , then we are forced to say that the Orioles' losing record and their lack of free-agent acquisitions aren't reasons for me to believe that the Orioles will perform badly this year. But this is surely not right; they are reasons for me to believe that the Orioles will perform badly this year, notwithstanding my failure to recognize that fact. And in particular, the evidentialist should want to say that. If I fail to believe that the Orioles will have a bad year – or, even worse, positively believe that they'll have a good year – I'm failing to respond to my evidence correctly, and thus failing to believe as I ought to; the evidentialist should want to capture that. But she can't if she's committed to restricting our reasons to the considerations that we *take* to be evidence.

Given this second, "too strong" problem, *Subjective Evidentialism-Reasons* seems implausible. But it did seem that the above argument from *Impossibility Claim-Precise* and *Constraint on Reasons* to *Subjective Evidentialism-Reasons* was valid. Do we, then, have a reductio of *Constraint on Reasons*?

Only on one interpretation of *Constraint on Reasons*, I think. Notice that *Constraint on Reasons* is not relativized to an agent. One interpretation of *Constraint on Reasons* would say that some consideration C constitutes a normative reason *for* A to Φ only if it's possible *for* A to Φ on the basis of C (holding fixed certain background mental states of A 's, such as her basic desires, or what she takes to be evidence for what).¹⁶ This would indeed seem to suggest that if I don't take some consideration to be evidence for p , then it can't be a reason for me to believe p , which yields the bad result in the Orioles case and many other cases with the same structure. But it might plausibly be suggested that this is not the right way to understand *Constraint on Reasons*. Rather, it might be said, this constraint should allow that C can constitute a normative reason for A to Φ so long as *someone* (with

¹⁶ This is close to the "reasons internalist" position of Williams (1981), who Shah (2006: 484) explicitly cites as having inspired his approach. It's not clear what strong, Williamsian reasons internalists should say about epistemic reasons, if their doctrine is supposed to apply in full generality.

the right psychology, which might include background mental states different to A's¹⁷) could Φ on the basis of C. Hence, the most plausible precisification of *Constraint on Reasons* is:

Constraint on Reasons-Precise. Some consideration C constitutes a normative reason for A to Φ only if it's possible for someone to Φ on the basis of C.

Constraint on Reasons-Precise doesn't rule out the Orioles' losing record from last year, plus their lack of free-agent acquisitions being a reason for me to believe that they will perform badly this year. It doesn't matter that I don't recognize the force of these considerations; the point is that someone who *did* would be able to believe that the Orioles will perform badly this year on this basis.

With this clarification made, though, we can see that the situation for the Kelly-Shah argument is even worse than it first appeared. We now have the following two premises:

Impossibility Claim-Precise. It is impossible for A to believe p on the basis of considerations that A takes to be evidentially irrelevant to whether p .

Constraint on Reasons-Precise. Some consideration C constitutes a normative reason for A to Φ only if it's possible for someone to Φ on the basis of C.

But what do *these* two claims jointly entail? Nothing (relevant), it seems. They don't entail *Evidentialism-Reasons*; they don't even entail *Subjective Evidentialism-Reasons*. (Thus even the argument for *Subjective Evidentialism-Reasons*, once its premises have been made plausible and precise, turns out to be invalid after all.) Moreover, they don't seem to rule out many (if any) would-be reasons for belief, since it seems that for any consideration that you take to be evidentially irrelevant to whether p , it's always going to be possible to imagine *someone* who does take it to be evidentially relevant to whether p , and is thus able to believe p on the basis of it.

To make this more vivid, return to the case where you're offered \$1,000,000 to believe p . As I noted above, by the lights of *Subjective Evidentialism-Reasons*, there's nothing to rule this offer of money out from counting as a normative reason for you to believe p if you *take* the offer of money to be evidence that p . But now we can see that, at least by the lights of the precisified constraint on reasons alone, there's nothing to rule the offer of money from counting as a normative reason for you to believe p even if you *don't* take the offer of money to be evidence that p . For it's still true that *someone* (specifically, someone who *did* take the fact that they've been offered \$1,000,000 to believe p to be evidence for p) could believe p on the basis of the fact that they've been offered \$1,000,000. And that's enough for this not to be ruled out by *Constraint on Reasons-Precise* even as a reason *for you* to believe p .

Although I seem to have shown that, when we make the premises and conclusion of the Kelly-Shah argument precise, it's not valid, one may have the sense that there's something fishy about the way I'm resisting it. Suppose for the sake of argument that the fact that you've been offered \$1,000,000 *does* constitute a (pragmatic, non-evidential) reason to believe p . I've been insisting that it *is* possible to believe p on the basis of this fact, but only by *mistaking* this fact for an *evidential* reason to believe p .

¹⁷ We could, if you like, think of this someone as a hypothetical version of A herself *if* her background mental states were to be changed in the relevant ways.

Isn't there something at least a bit odd about the idea that this fact *is* a pragmatic reason to believe p , but that one can only base one's belief in p on this fact by mistaking it for an evidential reason to believe p ?

I agree that there's something that calls out for explanation there, and I'll turn to this point later, in §4. For our purposes in this section, the question is whether there's a way to save something of the form of the Kelly-Shah argument itself in the vicinity here. I can see two strategies for claiming that there is.

So far I've been appealing to the undifferentiated notion of a "consideration", where the fact that you've been offered \$1,000,000 to believe p is the *same* consideration, whether you're thinking of it as evidence that p , or merely as something that shows believing p to be pragmatically valuable. One might wonder if there's a more fine-grained notion than that of a consideration, that will be sensitive to whether this consideration is serving as (or being thought of as?) evidence for p , or whether it is serving as (or being thought of as?) something that shows believing p to be pragmatically valuable. We'd then need to revise the *Constraint on Reasons* so that it requires it to be possible to base one's belief on this more fine-grained thing, in order for the corresponding reason to be present. Then, the thought would be that while it is possible to base your belief on the fine-grained entity (whatever it is) that corresponds to the fact that you've been offered \$1,000,000 *qua* (apparent?) evidence that p , it isn't possible to base your belief on the different fine-grained entity (whatever it is) that corresponds to the fact that you've been offered \$1,000,000 *qua* (apparent?) thing-that-shows-believing- p -to-be-pragmatically-valuable. And so the latter cannot constitute a normative reason.

While this is an intriguing strategy, a lot more has to be said to make it clear and to vindicate it. Most obviously, we're owed an account of the what the fine-grained entity is, exactly. Additionally, the view seems to be committed to saying that both motivating reasons (or bases) and normative reasons are constituted by this fine-grained entity, rather than by a coarser-grained one. This requires a significant rethinking of orthodoxy, since the standard view (which Kelly and Shah do not seem to depart from in the way they speak) is that reasons are *facts*, which are clearly too coarse-grained to do the job (the fact – that believing p will get you \$1,000,000 – is the same in both cases).¹⁸ Other prominent views, such as that reasons are propositions, or that they are mental states (likely, when it comes to epistemic reasons, *beliefs* in those propositions), seem likewise too coarse-grained to do the job. This rethinking of orthodoxy is a major task to be executed, and it seems *ad hoc* if it cannot be motivated independently of the need to save evidentialism.¹⁹

A different way of pursuing a proposal in the same broad spirit – this is the second strategy mentioned above – is to try to build the fine-grainedness not into our account of what a reason is, but rather into *Constraint on Reasons*.²⁰ So, one might try something like the following:

Constraint on Reasons-Revised. Some consideration C constitutes a normative reason of kind K for A to Φ only if it's possible for someone to Φ on the basis of C *under the guise of its being a*

¹⁸ For an influential defense of this view, see Dancy (2000: ch. 5).

¹⁹ Nathan Howard (ms.) is working out a view of reasons that may fit the bill, though he himself does not seem to think that non-evidential considerations aren't normative reasons; merely that we need fine-grained individuation in order to be able to classify them as "wrong-kind" reasons.

²⁰ I thank Maria Lasonen-Aarnio for this suggestion.

reason of kind K – that is, through an appreciation of the *distinctive kind of way* that C counts in favor of Φ -ing.

The thought here is that in order for the fact that you’ve offered me \$1,000,000 to believe p to constitute a genuine, practical reason for me to believe p , it has to be possible for me (or at least, someone) to believe p on the basis of this monetary offer *under the guise of a practical consideration* (rather than, say, by taking this fact to be evidence for p). That is – assuming for *reductio* that this consideration does practically count in favor of believing p – I (or someone) have to be able to believe p on the basis of the monetary offer *through appreciating the way in which in favors of favoring of believing p* – viz., practically. Put together with the premise that we can’t believe on the basis of practical considerations *under the guise of practical considerations*, this yields *Evidentialism-Reasons*.

I will be honest here: I think this is the best strategy for reviving the argument from a motivational constraint, and it is certainly worthy of consideration. That said, once again, this option is so far completely unexplored in the literature. In particular, I know of no-one who has argued for *Constraint on Reasons-Revised*, as opposed to the (weaker) *Constraint on Reasons-Precise*.

Here’s one challenge for this strategy. Most ordinary agents do not have sophisticated views, when they act or believe on the basis of some consideration, about exactly *how* this consideration counts in favor of the act or belief in question, or about what categories of reasons there are (viz. moral, prudential, instrumental, aesthetic, epistemic, and so on) and what the boundaries between them are, in a way that allows them to determine what kind of reason the consideration that they are acting or believing under the guise of is, such as to fix what kind of reason they are acting or believing under the guise of. Even philosophers find these questions very difficult to answer, let alone ordinary people. So a lot of the time it is going to be seriously indeterminate under what “guise” someone is acting under when they act on some consideration.

This is most obvious in the case of reasons for action, where it seems frequently indeterminate whether someone is acting on some consideration under the guise of a moral reason or under the guise of some other reason (prudential, instrumental, aesthetic, etc.). But it can also arise in the case of belief as well. Consider (a version of) the sort of case discussed by Stroud (2006) and others, where friendship impinges upon our doxastic practices. Suppose that Tamara is accused of a crime, and protests her innocence to her friend Smruti. Smruti then reports that she believes that Tamara is innocent “because she’s my friend” – not an unusual sort of thing for a person to say. It seems reasonable to take Smruti at her word when she says she believes Tamara because Tamara is her friend. But under what guise is Smruti conceiving of this consideration? She could be taking the fact that Tamara is her friend (someone she knows well, trusts, and so on) to be evidentially indicative of the fact that she wouldn’t lie to her; or she could be taking the fact that Tamara is her friend to be some kind of practical (moral? *sui generis* friendship-based?) reason to believe her. Or – and I suggest this is plausibly true in a wide range of real-life cases – it may simply be indeterminate whether Smruti is believing on the basis of her friendship under the guise of an evidential reason or under the guise of a practical reason. She may not have separated out, in her mind, the two kinds of (putative) ways that her friendship with Tamara might count in favor of believing Tamara – or considered in which way (or both) she takes it to so count in favor.

This makes it hard to definitively say whether Smruti counts as someone who can believe on the basis of her friendship with Tamara *under a practical guise*, and thus makes it hard to see whether *Constraint on Reasons-Revised* rules out friendship-based practical reasons for belief or not.²¹ Now, the proponent of *Constraint on Reasons-Revised* might insist that to settle whether it is possible for anyone to genuinely believe (or act) on some consideration C under the guise of a reason of some particular kind K, we should focus only on agents who have the distinction between the K-reasons and the non-K-reasons clearly in view, and see whether *they* can believe (or act) based on C under the guise of a K-reason. But such people are, in the main, sophisticated, professional philosophers. And it seems rather *ad hoc* – not to mention rather far from anything that motivated the *Constraint on Reasons* in the first instance – to insist that what kinds of reasons ordinary people can have ultimately turns only on how people quite unlike them, namely professional philosophers, are motivated.

While I have raised some initial challenges for the two strategies for saving the Kelly-Shah-style argument – the fine-grained view of reasons, and the appeal to different guises – I admit that these challenges are not entirely conclusive, and may ultimately be answerable. So I have to make do with the conditional conclusion that the Kelly-Shah-style argument can't be made to work *unless* one of these two strategies can be vindicated. It's in any case interesting if it turns out that the only way to preserve the argument is to embrace one of these views. In any case, I think it's clear that Kelly and Shah themselves didn't have either of these views in mind, and did not see the problem that requires them to adopt one of them. They simply weren't attending sufficiently to the difference between considerations that the evidence *actually* supports and considerations that one *takes* it to support, and the significance of this difference for their arguments.

3. Arguments from *ought* implies *can*

I next want to discuss another argument that is in the neighborhood of, but not identical to, the one we just considered – though the two are often not well-distinguished. I haven't found a clear statement of this argument as distinct from the first, but it seems to me to be “in the air” in discussions of evidentialism and pragmatism.²² The argument turns on an “ought implies can” principle. An initial rough statement of it might be as follows.

1. If there were non-evidential reasons for belief, then it would sometimes be the case that you ought to believe p when your evidence doesn't support p .
2. But you cannot believe p when your evidence doesn't support p . (impossibility-claim)

²¹ Remember that if *at least one person* can believe on the basis of some consideration C under a practical guise, then *Constraint on Reasons-Revised* does not rule it out as counting as a reason for *anyone*. This is a result of the (orthogonal) clarification that we issued in formulating *Constraint on Reasons-Precise*, which survives in the revised version. We could try formulating *Constraint on Reasons-Revised* so that it requires the actual individual for whom C is (putatively) a reason to be able to act on it under the guise of the kind of reason it is. But seems much too demanding. Those who lack the concepts required to distinguish between different kinds of reasons, or who have mistaken theories of the boundaries between those different kinds of reasons, can't act on those reasons *under the relevant guise*, but they don't thereby seem to lack those reasons.

²² Reisner (2009: 264-268) seems to be discussing such an argument, though he doesn't clearly distinguish it from the one we considered in the last section.

3. You ought to Φ only if you can Φ . (ought implies can)

So,

Evidentialism-Reasons. There are no non-evidential reasons for belief (all reasons for belief are evidential).

But, for reasons that are now familiar, premise (2) is wrong as it stands: what is true *at most* is that you cannot believe p when you don't *take* your evidence to support p . With premise (1) also adjusted to preserve the argument's validity, the result is:

1*. If there were non-evidential reasons for belief, then it would sometimes be the case that you ought to believe p when you don't take your evidence to support p .

2*. But you cannot believe p when you don't take your evidence to support p .

3. You ought to Φ only if you can Φ (ought implies can).

So,

Evidentialism-Reasons. There are no non-evidential reasons for belief (all reasons for belief are evidential).

The problem with this argument, though, is that it seems to overgeneralize. For even if there *aren't* non-evidential reasons for belief, it is sometimes the case that you ought to believe p even though you don't take your evidence to support p . This is so *by evidentialist lights* in cases where your evidence does decisively support p , but you don't recognize this. But given premises (2*) and (3), it seems that it will follow that it isn't true that you ought to believe p in such cases. Thus premises (2*) and (3) seem to threaten evidentialism itself.²³ Something in premises (2) and (3) seems to require toning down.

I suggest that the problem is this. (2*) is scope-ambiguous, being able to be interpreted either of the following ways:

2*-Narrow. When you don't take your evidence to support p , you cannot believe p .

2*-Wide. You cannot (believe p when you don't take your evidence to support p).

I think we should deny that (2*-Narrow) is true, at least on any reading of 'cannot' that vindicates (3) (i.e., ought implies can).²⁴ Suppose you don't take your evidence to support p . In the relevant sense of 'can' that vindicates 'ought implies can', you still *can* believe p – by (reasoned) change of mind about

²³ Reisner (2009: 266-267) raises a similar point.

²⁴ A semantic complication here is that on the standard semantics for modals (Kratzer 1981), the antecedents of conditionals act as restrictors for the semantic interpretation of their consequents. So this means that the antecedent 'when you don't take your evidence to support p ' restricts us to worlds in which you don't take your evidence to support p , such that the consequent – 'you cannot believe p ' comes out true just if, *quantifying just over those worlds*, none of them are ones in which you believe p . This means that the natural semantic interpretation of 2*-Narrow mimics the truth-conditions of 2*-Wide, and thus will be true. This is immaterial for my purposes, though. As you'll see below, if (2) is interpreted as (2*-Wide), the argument as a whole comes out invalid, and making it valid again requires an adjustment of premise (1) that makes it false. The same is true if (2) is interpreted as (2*-Narrow), and (2*-Narrow) is in turn interpreted in a way that makes its truth-conditions equivalent to (2*-Wide). So the only alternative is to interpret (2) as (2*-Narrow) where the conditional is stipulatively material, and not a natural language 'restrictor' conditional a la Kratzer. And *then*, it's false for the reasons that I'm suggesting.

whether the evidence supports p . That's why, in a case where your evidence *does* support p , we can still say that you ought to believe p , without any violation of (a plausible version of) ought implies can. The relevant sense of 'can' for ought implies can doesn't treat all your existing background psychological states as unalterably fixed, for this would be too restrictive. (Compare: you don't care about anyone except for yourself. In a weak sense that treats this not caring as unalterably fixed, this means that you *can't* act from altruistic motives. But it doesn't follow that it's false that you ought to act from altruistic motives. In the relevant sense, you *can* start caring about others, and thus start acting out of altruistic motives.)

That leaves (2*-Wide), which is more plausible (though still highly contestable). What it says is that you can't simultaneously (i) believe p and (ii) not take your evidence to support p . Suppose for the sake of argument that this is true. If we interpret (2) as (2*-Wide), we need to make a corresponding adjustment to (1) to preserve the argument's validity, giving us yet a third version of the argument:

1*-Wide. If there were non-evidential reasons for belief, then it would sometimes be the case that you ought to (believe p while simultaneously not taking your evidence to support p).

2*-Wide. But you cannot (believe p while simultaneously not taking your evidence to support p).

3. You ought to Φ only if you can Φ (ought implies can).

So,

Evidentialism-Reasons. There are no non-evidential reasons for belief (all reasons for belief are evidential).

But the anti-evidentialist who accepts (2*-Wide) and (3) should, and can, simply deny (1*-Wide). She should hold that when you have decisive non-evidential reasons to believe p , you should not only believe p , but also believe that your evidence supports p . After all, if (2*-Wide) is true, then believing that your evidence supports p will be a necessary means to succeeding in believing p .²⁵ Indeed, ever since Pascal, pragmatists about reasons for belief have counselled those who have pragmatic reason to believe things to do so *by* convincing themselves that there's sufficient evidence for these claims.²⁶ But if this is right, it won't be true that, in such cases, you ought to (believe p while simultaneously not taking your evidence to support p). Thus, (1*-Wide) will be false. Thus, ultimately, one way or another, the argument fails.

4. Arguments by inference to the best explanation

A third kind of argument from an impossibility-claim to evidentialism can be found in Jonathan Adler's book *Belief Own Ethics* (Adler 2002b: ch. 1). Though Adler doesn't quite explicitly present it this way, I think the argument can be presented as proceeding via inference to the best explanation. The idea is that it is a fact that we can't respond to non-evidential considerations, and that the best

²⁵ Of course, an anti-evidentialist who denies (2*-Wide) should not say this, since for them, believing that your evidence supports p won't be a necessary means to succeeding in believing p . But either way, one of the premises of the argument gets denied.

²⁶ Cf., e.g., Pascal (1670/1910: §233); Foley (1987: 216-222).

explanation of this fact is that we all tacitly recognize and accept the truth of *Evidentialism-Ought* (which, to recap, says that one ought to believe p only if one has adequate evidence for p).²⁷

If this style of argument works, it may also provide support for the claim that *Evidentialism-Ought* is a *constitutive* norm for belief. It is not completely clear what it is for a norm on some mental state to be constitutive of that mental state, as opposed to being a genuine, necessary, norm on that mental state that is nevertheless not constitutive of it. But here is an initially appealing thought: for many norms N on some mental state kind K , it doesn't follow merely from N 's being a norm on states of kind K that everyone tacitly recognizes and accepts it as a norm on states of kind K . So maybe what distinguishes the *constitutive* norms on states of kind K are that they are ones that anyone who counts as having a mental state of kind K – or, perhaps, anyone who has the *concept* of mental state kind K , such that they're capable of thinking of their own mental states as falling under kind K ²⁸ – must tacitly recognize and accept, and thus be to at least some extent disposed to comply with. If this is the right way of drawing the distinction between constitutive and non-constitutive norms is correct, then anyone who is committed a claim of the form:

Subjection Constitution. It's constitutive of being in a state of kind K that one's state is *subject* to norm N .

is also committed, given that they think that N is a *constitutive* norm for S , be committed to:

Compliance Constitution. It's constitutive of being in a state of kind K (or, perhaps, of conceptualizing one's state as being of kind K) that one is at least to some extent disposed to regulate one's state in compliance with N .²⁹

Given this account of the difference between constitutive and non-constitutive norms, it makes sense for someone who accepts the Adlerian argument to also accept that *Evidentialism-Ought* is a constitutive norm. For, according to Adler, the best explanation of the fact that we can't respond to non-evidential considerations is that we all (or, at least, those of us who have the concept of belief) tacitly recognize and accept the truth of *Evidentialism-Ought*. And if that's so, then by this account of constitutivity, *Evidentialism-Ought* is a constitutive norm: it will be constitutive of being in the state of belief both that this state is subject to *Evidentialism-Ought*, and that (perhaps just insofar as one conceives of the state as a belief) one is disposed to regulate it in compliance with *Evidentialism-Ought*.³⁰

²⁷ See esp. Adler 2002b: 25, 29.

²⁸ Depending on whether the constitutive claim is metaphysical or conceptual.

²⁹ See Barranco Lopez (ms.) for a similar argument (and a similar way of drawing the distinction between constitutive and non-constitutive norms). Subjection Constitution claims are not always well-distinguished from Compliance Constitution claims, but in fact many constitutivists do seem to endorse the latter as well as the former. For example, Wedgwood (2017: 4) writes that “the norms of rationality are in a way *constitutive* of the various different types of mental state. According to the version of this claim that I shall advocate, if one is even to be interpretable as having mental states of the relevant types at all, one must have at least some disposition to *conform* to these norms” (my italics).

³⁰ This fits with what Adler says, as he repeatedly stresses that evidentialism is built into the very concept of belief (e.g. Adler 2002b: 1, 10, 25, 32) – which he at times seems to equate with its being built into the *nature* of belief (e.g. Adler 2002b: 4).

As with the previous arguments, there is a question about how to interpret the claim that we can't respond to non-evidential considerations, as it appears as a premise in the argument by inference to the best explanation: as referring to considerations that aren't *actually* evidence, or as referring to considerations that we don't *take* to be evidence. And as we've already seen in §1, Adler is fairly uncareful about moving between these two sorts of claims. Nevertheless, at least some of the time, he presents the datum to be explained as being the fact that we can't respond to considerations we don't *take* to be evidence, or that we can't defy *our own judgments* about what the evidence supports.³¹ And as we've seen, this is how we need to interpret the datum if it's to be at all plausible. (Indeed, we may want to weaken it further, to say, for example, that we're merely *strongly disposed* not to defy our own judgments about what the evidence supports.)

Now, we might hold out some hope that, unlike the arguments of the last two sections, the argument by inference to the best explanation still works well even when this is taken account of. My recognition and acceptance of the truth of evidentialism sounds like a good candidate for explaining why I can't respond to considerations that I don't *take* to be good evidence, since I will (normally) *take* responding to them to involve having to violate evidentialism, a norm that I accept and am disposed to comply with. This evidentialist explanation reissues the challenge that I mentioned toward the end of §2: if there are, contrary to this explanation, non-evidential reasons for belief, what *can* explain why we only respond to them by mistaking them for evidential considerations?

However, once we make clear that what we need to explain is not a disposition to believe what the evidence actually supports, but a disposition to believe what one *takes* the evidence supports, I think a competing explanation comes into view. As a number of philosophers have suggested, there is a coherence requirement on doxastic attitudes, distinct from evidentialism itself, that requires us to line up our beliefs with our *judgments* about what the evidence supports:

Inter-level coherence (ILC). Rationality requires that you do not simultaneously:

- (i) {Believe that a doxastic attitude D is inadequately supported by the evidence, have D}; *or*
- (ii) {Believe that a doxastic attitude D is decisively supported by the evidence, fail to have D}³²

Failing to distinguish *Inter-Level Coherence* from *Evidentialism-Ought* is a mistake very similar to failing to distinguish what one *takes* one's evidence to support from what is *actually* supports, since *Inter-Level Coherence* demands conformity of one's beliefs with what one *takes* one's evidence to support rather than with what it actually supports.³³ Despite their distinctness, though, *Evidentialism-Ought* and *Inter-*

³¹ See e.g. Adler 2002b: 15, 26, 29, 67.

³² I've defended this principle elsewhere (Worsnip 2018a). See also, e.g., Horowitz (2014), Greco (2014) and Titelbaum (2015), among others. Not everyone who endorses the claim that the states forbidden by ILC are always irrational is clearly thinking of ILC as a *sui generis* coherence requirement independent of evidentialist requirements like *Evidentialism-Ought*. Some just present the point as being that one's evidence could never support both of the states in (i) simultaneously, nor both of the states in (ii). (Cf., e.g., Horowitz 2014, and more explicitly, Kiesewetter 2017: 248-254.) I argue for thinking of it as a *sui generis* requirement in Worsnip 2018a.

³³ Arguably, Adler is an example of someone who fails to do this; likewise, Feldman (2005) seems to move between the two norms in a way that isn't fully clear.

Level Coherence are at least *prima facie* consistent, due to the latter's "wide-scope" character. Suppose you falsely judge that your evidence supports believing *p*. In that case, *Evidentialism-Ought* requires you *not* to believe *p*. But *Inter-Level Coherence* doesn't categorically demand that you do believe *p*, since you could also satisfy it by giving up your belief that your evidence supports believing *p*.³⁴

As far as I can tell, one could mimic Adler's explanation of why we can't respond to considerations that we take to be non-evidential in belief-formation using *Inter-Level Coherence* instead of *Evidentialism-Ought*, and not much is lost from the explanation. The suggestion would be that we all tacitly recognize and accept the irrationality (or incoherence) of simultaneously holding, for example, that the evidence doesn't support *p*, while also believing *p* and thus holding that *p* is true – and that this is why we have difficulty responding to pragmatic reasons in belief-formation.³⁵ At least *if* one finds the style of explanation in terms of *Evidentialism-Ought* satisfying, I don't see why one shouldn't find the explanation in term of *Inter-Level Coherence* equally satisfying (pending some worries to be considered in the next section).

If that's so, the argument for *Evidentialism-Ought* by inference to the best explanation is unsound. For an argument to a conclusion by inference to the best explanation to work, the conclusion has to be the *single* best explanation of the relevant phenomena or data. If there are other, equally good, explanations available, then it seems like the argument at best provides an argument for the disjunction of the good explanations. So, given that there is a rival explanation of our inability to respond to pragmatic reasons, involving *Inter-Level Coherence*, that seems equally good to the one involving *Evidentialism-Ought*, what we have is at best an argument for the disjunction of *Evidentialism-Ought* and *Inter-Level Coherence*.

This also raises the intriguing possibility that it is *Inter-Level Coherence*, rather than *Evidentialism-Ought*, that is a *constitutive* norm for belief. Indeed, if the mark of a constitutive norm on belief is that any believer (or, perhaps, any believer with the concept of belief) must be disposed to comply with it, then it seems that the case for the constitutivity of *Inter-Level Coherence* is actually better than the case for the constitutivity of *Evidentialism-Ought*. If we are disposed to believe what we *take* our evidence to support, this more directly and obviously constitutes a disposition to comply with *Inter-Level Coherence* than a disposition to comply with *Evidentialism-Ought*. It's a disposition to do what we *think* will bring us into compliance with *Evidentialism-Ought*, but given the relationship between *Evidentialism-Ought* and *Inter-Level Coherence*, that is very close to just restating the disposition to comply with *Inter-Level Coherence*. There seems to be a more direct sense, then, in which we're disposed to comply with *Inter-Level Coherence* than in which we're disposed to comply with *Evidentialism-Ought*. This bolsters the case for thinking of the former, rather than the latter, as constitutive.³⁶

³⁴ Cases, if there are any, where you falsely *but rationally* believe that your evidence supports believing *p* are more complex; see Worsnip 2018a for discussion.

³⁵ Indeed, something in the neighborhood this explanation of the difficulty of responding to pragmatic reasons in belief-formation is pursued by Winters (1979) and Setiya (2008).

³⁶ Indeed, I've suggested elsewhere (Worsnip 2018b) that quite generally, it's coherence requirements – also known as requirements of "structural rationality" – that are constitutive of mental states, rather than "substantive" norms like evidentialism. It's a distinctive feature of coherence requirements, as opposed to substantive norms, that the former are hard to transparently violate in a way that the latter are not.

5. Doubts about an ILC-based explanation, considered

The primary purpose of this paper isn't to develop an explanation of our inability to believe for non-evidential reasons in terms of *Inter-Level Coherence*; it is rather to put pressure on various arguments for evidentialism from our inability to believe for non-evidential reasons. However, for the response that I gave to the argument by inference to the best explanation in the last section to work, it does need to be the case that the *ILC*-based explanation is at least as good as the evidentialist explanation (and that the two are independent of one another). This section considers three different doubts that one might have about this.

a) *Is ILC derivative on evidentialism?*

A first worry is that the appeal of *Inter-Level Coherence* is just derivative on the appeal of evidentialism. Notice that *Inter-Level Coherence* specifically bans defying your judgments about what the *evidence* supports, not defying your judgments about what you all-things-considered ought to believe. One might think that this is only attractive insofar as we already accept that evidence is determinative of what one all-things-considered ought to believe, which is a kind of evidentialism.³⁷ If *Inter-Level Coherence* is itself explained by the truth of evidentialism, then an explanation of the difficulty of responding to pragmatic reasons in terms of the former ultimately bottoms out in terms of the latter, and so is not a rival to an evidentialist explanation after all, thus putting Adler's argument back on a more secure footing.

However, I do not agree that *Inter-Level Coherence* should be attractive only insofar as one is antecedently committed to evidentialism. I think of *Inter-Level Coherence* as being part of a family of *sui generis* coherence requirements – ranging across all attitudinal mental states including belief, desire, intention, hope, and so on – that are not necessarily derivative on or explained by substantive norms. The relevant claim is that there is some kind of internal incoherence in (for example) believing that one lacks adequate evidence for *p*, but nevertheless believing *p* all the same. What explains why it's incoherent to believe *p* while believing that one lacks *evidence* for *p* specifically? (As opposed to, say, it's only being incoherent to believe *p* while believing that one lacks *adequate reason of some kind* for believing *p*?) In my view, the answer is found in the conceptual connection between evidence, truth and belief. To believe *p* is to believe that *p* is true, but evidential considerations with respect to *p* just are those that bear on whether *p* is true. Contra evidentialists, I don't think this proves that only evidence can be a reason for belief, but I do think it explains why it's internally incoherent to believe *p* while believing that one lacks adequate evidence for *p* (or to fail to believe *p* while believing that one has decisive evidence for *p*).

As I stressed at the outset, I myself am not taking a stance in this paper on whether evidentialism is true, merely contending that arguments from our inability to respond to non-evidential reasons don't show evidentialism to be true. Nevertheless, suppose that evidentialism *isn't* true – and that there are substantial, pragmatic reasons for belief, such that you sometimes ought to believe, all-things-considered, against your evidence – but *Inter-Level Coherence* is true, so that you'd be irrational

³⁷ Adler (2002b: 9-12) assumes that the incoherence of states that violate *ILC* is evidence for evidentialism.

to believe against what you *take* to be your evidence. You can believe what you all-things-considered ought to believe without irrationality if you have a way of manipulating your beliefs about what your evidence supports. Since *Inter-Level Coherence* is a synchronic requirement, all that's required to respect it is that there's no *single* point in time when you (for example) believe *p* while believing that your evidence doesn't adequately support *p*.

Again, this fits quite comfortably with the standard advice that pragmatists give about how to respond to pragmatic reasons. They counsel doing so *via*, for example, spending time in environments that will alter your evidential standards (for example, going to church if you want to believe in God), or engaging in forms of selective evidence-gathering.³⁸ So pragmatists are already counseling us to respond to pragmatic reasons in ways that respect *Inter-Level Coherence*. (Indeed, this isn't an accident if we're all unavoidably disposed to satisfy ILC, since these may be the only ways of effectively responding to pragmatic reasons.³⁹) Thus, I don't think that adding *Inter-Level Coherence* to a pragmatist picture is unnatural or disjointed.

Interestingly, some *evidentialists* also counsel us to respond to pragmatic considerations by changing our evidential standards or engaging in selective evidence-gathering. (Typically, and unlike the pragmatist, they say that the pragmatic value of believing *p* doesn't give you a reason to *believe p*, but *only* a reason to "get yourself" to believe *p*.⁴⁰) Some may even think that doing this (often) enables us to respond to pragmatic considerations without irrationality. But this claim is a mistake, by evidentialist lights, one that again seems to reflect a failure to distinguish what the evidence supports and what one takes it to support. Convincing yourself of a new set of evidential standards, according to which your evidence supports believing *p*, is not the same thing as making it the case that your evidence *actually* supports believing *p*. Even if we focus only on selective evidence-gathering, it is very questionable whether one can actually make it the case that one's evidence supports *p* by engaging in an evidence-gathering process that is deliberately designed to filter out any counter-evidence against *p* – since the probative value of the apparent evidence for *p* that you gather is significantly diminished, if not eliminated entirely, if you know that you received a "biased sample" that included the evidence for *p* but not the evidence against.⁴¹

Given this, the evidentialist who counsels you to "get yourself" to believe *p* in response to pragmatic considerations is committed to counseling you to do something that will result in your violation of *Evidentialism-Ought*. It seems that the evidentialist has two options to avoid outright

³⁸ See also Kunda (1990), who suggests that the psychological research shows that, a descriptive matter, practical desires tend to influence beliefs via causing us to selective filtering of and attention to evidence.

³⁹ Of course, there might be *some* situations where it's pragmatically valuable to believe *p*, but you don't have a way of manipulating your beliefs about whether your evidence supports *p*. In those situations there isn't a clear route that you can deliberately take to believing *p*, rationally or otherwise. But that's life, or so it seems to me. The 'ought implies can' considerations discussed in §3 may reappear here, but even if they do, they establish at most that there are *some* situations where pragmatic considerations fail to make it the case that you ought to believe *p*, which falls far short of evidentialism.

⁴⁰ Cf., e.g., Mills (1998); Kelly (2002: 171); Shah (2006: 493-6).

⁴¹ See Titelbaum 2010 and (especially) Salow 2018. Here's an example, adapted from one used by Roger White (2010) in a different context. Suppose you go to a party that is being attended by experts about economic forecasting. You go around asking these experts whether they expect inflation to go up this year, and to your surprise they all agree that it will. But then you discover that the host of the party asked each attendee at the door whether they expect inflation to go up. Those who said 'yes' were admitted to the party, while those who answered 'no' were quietly led down to the basement and shot. Once you know your evidence has been selectively filtered in this way, it seems to lose almost all of its probative value.

contradiction here. The first option is to say that you should do something that will result in your violating *Evidentialism-Ought*, but you shouldn't violate *Evidentialism-Ought*. This seems very awkward, though: it is close to the advice 'get yourself to believe p , but don't believe p !'. Moreover, surely to the extent that you shouldn't violate *Evidentialism-Ought*, you shouldn't (or at least have a fairly strong reason not to) do things that will result in your violating it. If evidentialism is a norm that we shouldn't care at all about violating, the sense in which we should respect it seems somewhat anemic. The second option is to say that you *epistemically* ought to satisfy *Evidentialism-Ought*, but pragmatically ought to violate it. But saying that is very close to conceding that *Evidentialism-Ought* is true only on a stipulatively epistemic reading of 'ought', and not on an all-things-considered reading. This stance is completely consistent with allowing for pragmatic reasons for belief (and thus, the falsity of *Evidentialism-Reasons*). So neither of these options is comfortable.

Thus, if anything, I think it's the evidentialist – at least, the evidentialist who recognizes pragmatic reasons to *get ourselves* to believe things – that ends up with the more disjointed account, and not the pragmatist who supplements her view with coherence constraints like *Inter-Level Coherence*.

b) *Is ILC a "requirement" that is impossible to violate?*

A second worry is this: I've claimed that *Inter-Level Coherence* is a requirement of (structural) rationality. But at least some of the time (e.g., in premise (2) of the argument from ought implies can), the impossibility-datum that we were trying to explain has been characterized as the claim that it's *impossible* to defy one's own judgments about what the evidence supports. But if that's so, it seems to be impossible to violate *Inter-Level Coherence*. Not much of a requirement, if we can't violate it.

The question, which I've been bracketing for the sake of argument until now, is whether that version of the impossibility-datum is too strong. I think it is. I *do* think there's something importantly puzzling and hard to interpret about someone who claims to (e.g.) believe p but also take the evidence not to support p .⁴² But we can ultimately make sense of such a person if we don't assume that their psychological states are totally transparent to them. What is not possible is someone who, in full transparency, *recognizes* that they are simultaneously in the state of believing p and that of believing that the evidence doesn't support p , and has no disposition to revise their beliefs.⁴³ Any state that can withstand *recognition* that it's being held in defiance of the agent's own judgment about the evidence just doesn't count as a belief: it is some other state. This is a corollary of ILC's being a *constitutive* norm in the sense that the claim that I called "Compliance Constitutivism" is true of it: to count as having

⁴² Though there has been a considerable literature on whether "epistemic akrasia" (roughly, the violation of ILC) is impossible, with several philosophers arguing for an affirmative answer (e.g. Hurley 1989: 130-5, 159-70; Adler 2002a), the recent normative debate about "epistemic akrasia" largely ignores this literature, simply taking the possibility of epistemic akrasia for granted (and often not even really engaging the weaker idea that there's something very (psychologically) puzzling about it). It's interesting that this contrasts so sharply with the near-orthodoxy of the claim that it's impossible to directly and consciously respond to pragmatic reasons, given that this seems an equally bold claim, and one that may be quite closely related.

⁴³ See Worsnip 2018b for more. Ultimately, Adler himself qualifies his own denial of the possibility of epistemic akrasia with a similar "full awareness" condition (Adler 2002b: 29). But – again failing to distinguish what our evidence actually support and what we take it to support – he seems to move back and forth between this claim and the claim that it's impossible to violate *evidentialism* under conditions of full transparency. See n. 7 above on why I don't accept *that* claim.

the states that it governs, an agent must have some disposition to satisfy it. I'm suggesting that this disposition is a disposition to satisfy it *when conditions of full transparency are met*.

This still allows for violation of *Inter-Level Coherence*, since our states often aren't transparent to us. At the same time, it still allows for an explanation of why it's typically impossible for us to respond to considerations that we take to be evidentially irrelevant in belief-formation, since doing so would, in the paradigm cases, transparently involve our coming to believe something for which we judge ourselves to have insufficient evidence.

c) *Does the ILC-based explanation undergenerate?*

But the 'typically' qualifier in this last contention raises a question. Aren't there *some* cases where responding to (what one recognizes to be) pragmatic considerations would not involve violating *Inter-Level Coherence*? And so doesn't an explanation of the impossibility of responding to pragmatic considerations in terms of our disposition to satisfy *Inter-Level Coherence* fail to cover such cases, thus undergenerating?

Here are some cases in which responding to pragmatic considerations would *not* seem to involve a violation of *Inter-Level Coherence*:

1. Cases where one takes one's evidence to be permissive between two options (e.g. believing *p* and suspending judgment about whether *p*);
2. Cases where one has no firm view about whether one's evidence permits some doxastic attitude;⁴⁴
3. Cases where pragmatic considerations help to fix the threshold for *sufficient* (or *decisive*) evidence for believing;⁴⁵
4. Cases where one's believing *p* would itself rationalize a belief that one's evidence supports *p*,
 - a. Cases where the fact that one believes *p* provides apparent evidence that *p* (e.g., when believing that one will succeed makes it more likely that one will succeed).
 - b. Cases where whether there is strong evidence that *p* depends on whether *p* is true (e.g., arguably the belief that there is an external world; see Worsnip 2019).

However, in fact I think it is far from obvious that one can't respond to pragmatic considerations in these cases. The paradigm cases where one cannot respond to pragmatic considerations – cases like the \$1,000,000 offer to believe that the number of stars in the sky is even – are cases where one comes in with a fixed view that the evidence determine does *not* support believing the proposition in the

⁴⁴ Some (e.g. Adler 2002b: 311 n. 6) hold that it's also irrational to believe *p* in the absence of a judgment that one's evidence supports *p*. I don't think this is right, since it threatens to require us to have an infinite iteration of higher-order beliefs. Others (Horowitz 2014: 724; Feldman 2005: 118 n. 6; Huemer 2011: 1), more cautiously, think that it's irrational to believe *p* while suspending judgment about whether the evidence supports believing *p*. I still think this is not right. Suppose you suspend judgment about whether the evidence uniquely supports believing *p*, or uniquely supports suspending about whether *p*. If it were irrational to combine this with believing *p*, by parity of reasoning, it would also be irrational to combine it with suspending judgment about whether *p*. But then you'd be irrational *whatever* attitude you took in such a case.

⁴⁵ See, e.g., Owens (2000: ch. 2).

question (thus clearly not falling into any of the above categories). Moreover, they're cases where it's pretty obvious that this is so, and where any evidential standards that one could adopt that would license belief are so far away from one's current standards that they would constitute a drastic revision of one's standards and belief-forming dispositions.

By contrast, the best cases for those who think we sometimes *can* respond to pragmatic considerations are cases involving beliefs in domains where the epistemic standards for what counts as good evidence are much murkier and more contestable – moral and religious cases, in particular. It's quite plausible that those who do claim to be able to respond to pragmatic considerations in such cases are thinking of such cases as falling into one of the above categories: they're not sure what to think the evidence supports, or they think the evidence is permissive (i.e. leaves things open), for example. If we think that responding to pragmatic considerations is possible in cases where it does not require transparently violating *Inter-Level Coherence*, but not possible in other cases, we get a nice explanation of the difference between the cases where responding to pragmatic considerations seems nigh-on impossible and those where it seems closer to being possible. Indeed, it would then be a positive advantage of the explanation in terms of *Inter-Level Coherence* over the explanation in terms of a substantive evidentialist norm that the former can accommodate the possibility of responding to pragmatic considerations in certain special cases where this seems possible.⁴⁶

This is a lot to ask you to accept, though – arguing for this in detail would require a paper of its own – and I don't want to put too much weight on it. So let me point out, as a more defensive observation, that to the extent that you do think that the explanation involving *Inter-Level Coherence* undergenerates for explaining why we can't respond to pragmatic considerations (when we can't), it seems to me that an explanation involving evidentialism is likely to undergenerate too. For, recall that the explanation involving evidentialism *doesn't* go by saying that one can't violate evidentialism, or even that one can't violate evidentialism without some kind of failure of transparency to oneself. That, once again, would be to miss the difference between what one's evidence actually supports and what one takes it to support. Rather, it goes by saying that one can't *take oneself* to be violating evidentialism (without some failure of transparency). And that is *extremely* similar to the claim that one cannot violate *Inter-Level Coherence* (without some failure of transparency). So, to the extent that there are problems for the explanation in terms of *Inter-Level Coherence* undergenerating, they will likely apply to the explanation in terms of evidentialism as well. Maybe both explanations are hopeless; that would still block the argument by inference to the best explanation. My suggestion has just been that since the explanation in terms of *Inter-Level Coherence* is at least on a par with the explanation in terms of evidentialism, we can't infer to the latter by inference to the best explanation.

6. Conclusion

⁴⁶ There is a different kind of case in which responding to pragmatic considerations seems not to involve violating ILC, but where I would not say that it is thereby possible. Suppose you do believe yourself to have adequate evidence for *p*. Nevertheless, it seems not possible to believe *p* on the basis of some other consideration that you take to be evidentially irrelevant as to whether *p*. This suggests that we need to embrace a “basing” variant of ILC that forbids believing *p* *on the basis of some consideration q* while taking *q* to be evidentially irrelevant to whether *p*. I think this can also be rationalized in the same sort of way of the original ILC, appealing to the conceptual connection between evidence, truth and belief. Thanks to Bob Beddor both for raising the problem and suggesting the solution.

Where does all of this leave us? As I've said, in a sense this paper has been very negative, just trying to put pressure on a certain family of arguments for evidentialism rather than to directly argue against evidentialism itself. My own sympathies on these matters fluctuate, but in case you're interested, I'll lay my cards on the table. As I hinted earlier, I'm happy to accept a version of *Evidentialism-Ought* where the 'ought' is stipulatively epistemic, and not all-things-considered. I also think that there's a distinctive kind of evaluation, epistemic evaluation, that is concerned, primarily, with how well our beliefs fit the evidence. But I've never seen a convincing argument to suggest that we can't also evaluate beliefs in other terms, or that it can't be sensible to talk of practical reasons for belief. Moreover, there are positive reasons – that I touched on at the end of §5a – to be suspicious of this view. It's hard to deny that practical considerations could give us reasons to *get ourselves* to believe things, but I think the view that they give us reasons to get ourselves to believe these things, but not to believe them, is more awkward than most defenders of evidentialism have acknowledged. Thus, I'm inclined to reject both *Evidentialism-Reasons* and any version of *Evidentialism-Ought* that uses the all-things-considered ought.

At the same time, I think that to reject *Inter-Level Coherence* along with evidentialism would be to throw the baby out with the bathwater. As I've tried to argue, there are reasons to accept *Inter-Level Coherence* even if one rejects evidentialism. I think of *Inter-Level Coherence* as part of a broader family of coherence requirements on mental states (or, requirements of structural rationality), that are their own distinctive normative category, neither distinctively epistemic nor distinctively practical. I think that both the evidentialist and the anti-evidentialist should be willing to make room for them.

I also hope that, independently of the debate about evidentialism itself, this paper has been a case study of why it is important to be careful about the distinction between what the evidence supports and what one takes it to support (and the corresponding distinction between substantive norms like evidentialism and structural requirements like ILC). My only hypothesis about why this distinction gets elided is that people (both philosophers, and others) sometimes have a tendency to think of what the evidence supports as being obvious or luminous, such that if you just focus on the evidence carefully and conscientiously, what the evidence supports and what you take it to support will be the same thing. This isn't helped by a narrow focus on cases where what the evidence supports *is* pretty obvious. But the fact that some consideration evidentially supports some doxastic state is a normative fact,⁴⁷ and there are no general grounds for supposing that the normative facts about which considerations support which doxastic states (and how strongly) are any more obvious than the normative facts about which considerations support which actions (and how strongly). Though there are some cases where it's obvious what our reasons support doing, no-one would treat figuring out what our reasons support doing, *in general*, as an obvious matter. They shouldn't do so with respect to what our (evidential) reasons support believing, either.

References

Adler, J. (2002a). "Akratic Believing?," *Philosophical Studies*, 110/1: 1-27.

----- (2002b). *Belief's Own Ethics*. Cambridge, MA: MIT Press.

Barranco Lopez, A. (ms.). "How to Understand Normativism about the Nature of Belief."

⁴⁷ Cf. also Kelly 2006.

- Boghossian, P. (2014). "What is Inference?," *Philosophical Studies*, 169/1: 1-18.
- Conce, E. & Feldman, R. (1985). "Evidentialism," *Philosophical Studies*, 48/1: 15-34.
- Dancy, J. (2000). *Practical Reality*. Oxford: Oxford University Press.
- Feldman, R. (2005). "Respecting the Evidence," *Philosophical Perspectives*, 19: 95-119.
- Foley, R. (1987). *The Theory of Epistemic Rationality*. Cambridge, MA: Harvard University Press.
- Flowerree, A. (2017). "Agency of Belief and Intention," *Synthese*, 194/8: 2763-2784.
- Gendler, T.S. (2008). "Alief in Action (and Reaction)," *Mind & Language*, 23/5: 552-85.
- Greco, D. (2014). "A Puzzle about Epistemic Akrasia," *Philosophical Studies*, 167/2: 201-219.
- Horowitz, S. (2014). "Epistemic Akrasia," *Noûs*, 48/4: 718-744.
- Howard, N. (ms.). "Factualists Can't Solve the Wrong Kind of Reason Problem."
- Huemer, M. (2011). "The Puzzle of Metacoherence," *Philosophy and Phenomenological Research*, 82/1: 1-21.
- Hurley, S.L. (1989). *Natural Reasons*. Oxford: Oxford University Press.
- Kelly, T. (2002). "The Rationality of Belief and Some Other Propositional Attitudes," *Philosophical Studies*, 110/2: 163-196.
- Kiesewetter, B. (2017). *The Normativity of Rationality*. Oxford: Oxford University Press.
- Kratzer, A. (1981). "The Notional Category of Modality," in Eikmeyer & Rieser (eds.), *Words, Worlds, and Contexts*. Berlin: de Gruyter.
- Kunda, Z. (1990). "The Case for Motivated Reasoning," *Psychological Bulletin*, 108/3: 480-498.
- Leary, S. (2017). "In Defense of Practical Reasons for Belief," *Australasian Journal of Philosophy*, 95/3: 529-542.
- Mills, E. (1998). "The Unity of Justification," *Philosophy and Phenomenological Research*, 58/1: 27-50.
- Nolfi, K. (2018). "Why Only Evidential Considerations Can Justify Belief," in McHugh, Way & Whiting, *Normativity: Epistemic and Practical*. Oxford: Oxford University Press.
- Noordhof, P. (2004). "Believe What You Want," *Proceedings of the Aristotelian Society*, 101/1: 247-265.
- Owens, D. (2000). *Reason Without Freedom*. Abingdon: Routledge.
- Pascal, B. (1670/1910). *Pensées*. London: Dent.
- Reisner, A. (2009). "The Possibility of Pragmatic Reasons for Belief and the Wrong Kind of Reasons Problem," *Philosophical Studies*, 145/2: 257-272.
- Rinard, S. (2015). "Against the New Evidentialists," *Philosophical Issues*, 25: 208-223.
- Salow, B. (2018). "The Externalist's Guide to Fishing for Compliments," *Mind*, 127/507: 691-728.
- Setiya, K. (2008). "Believing at Will," *Midwest Studies in Philosophy*, 32: 36-52.
- Shah, N. (2003). "How Truth Governs Belief," *Philosophical Review*, 112/4: 447-482.
- (2006). "A New Argument for Evidentialism," *Philosophical Quarterly*, 56/225: 481-498.
- Steglich-Petersen, A. (2008). "Does Doxastic Transparency Support Evidentialism?," *dialectica*, 62/4: 541-547.
- Stroud, S. (2006). "Epistemic Partiality in Friendship," *Ethics*, 116/3: 498-524.
- Titelbaum, M.G. (2010). "Tell Me You Love Me: Bootstrapping, Externalism, and No-Lose Epistemology," *Philosophical Studies*, 149/1: 119-134.
- (2015). "Rationality's Fixed Point (or: In Defense of Right Reason)," *Oxford Studies in Epistemology*, 5: 253-294.

- Wedgwood, R. (2017). *The Value of Rationality*. Oxford: Oxford University Press.
- White, R. (2010). "You Just Believe That Because..." *Philosophical Perspectives*, 24: 573-615.
- Williams, B. (1973). "Deciding to Believe," in his *Problems of the Self*. Cambridge, UK: Cambridge University Press.
- (1981). "Internal and External Reasons," in his *Moral Luck*. Cambridge, UK: Cambridge University Press.
- Winters, B. (1979). "Believing at Will," *Journal of Philosophy*, 76/5: 243-256.
- Worsnip, A. (2018a). "The Conflict of Evidence and Coherence," *Philosophy and Phenomenological Research*, 96/1: 3-44.
- (2018b). "What is (In)coherence?," *Oxford Studies in Metaethics*, 13: 184-206.
- (2019). "What to Believe About Your Belief That You're in the Good Case," *Oxford Studies in Epistemology*, 6: 206-233.