Philosophy and
Phenomenological Research

Philosophy and Phenomenological Research
doi: 10.1111/phpr.12246

The Conflict of Evidence and Coherence

ALEX WORSNIP
New York University

For many epistemologists, and for many philosophers more broadly, it is axiomatic that rationality requires you to take the doxastic attitudes that your evidence supports. Yet there is also another current in our talk about rationality. On this usage, rationality is a matter of the right kind of coherence between one's mental attitudes. Surprisingly little work in epistemology is explicitly devoted to answering the question of how these two currents of talk are related. But many implicitly assume that evidence-responsiveness guarantees coherence, so that the rational impermissibility of incoherence will just fall out of the putative requirement to take the attitudes that one's evidence supports, and so that coherence requirements do not need to be theorized in their own right, apart from evidential reasons. In this paper, I argue that this is a mistake, since coherence and evidence-responsiveness can in fact come into conflict. More specifically, I argue that in cases of misleading higher-order evidence, there can be a conflict between believing what one's evidence supports and satisfying a requirement that I call "inter-level coherence". This illustrates why coherence requirements and evidential reasons must be separated and theorized separately.

For many epistemologists, and for many philosophers more broadly, it is axiomatic that rationality requires you to take the doxastic attitudes that your evidence supports.[1] Some may even take this fact to be a definitional truth, given the way they use the term 'rationality'. Yet there is also another current in our talk about rationality. On this usage, rationality is a matter of the right kind of *coherence* between one's mental attitudes. So, rationality requires one not to combine attitudes in various ways. It forbids one from simultaneously believing some proposition p and believing its negation; it forbids one from simultaneously believing p whilst assigning p some vanishingly small credence; it forbids one from simultaneously believing p whilst oneself judging one's reasons for believing p to be no good.

[1] Timothy Williamson (2000: 164) and Thomas Kelly (2006: §2) both call it a "platitude". See also, amongst many others, Conee & Feldman (1985); Owens (2000); Kelly (2002); White (2007); Fantl & McGrath (2009); Huemer (2011); Smithies (2012); Greco (2014).

How are these two currents of talk related? At least until very recently indeed,[2] very little work in epistemology has been explicitly devoted to answering this question. But various implicit answers have been given.

One answer effectively tries to reduce evidence-responsiveness to a kind of coherence. This is, I take it, more or less the idea behind the project of the "coherence theory of justification",[3] and well as that of some forms of subjective Bayesianism.[4] On this view, for a belief to be supported by the evidence just is for it to cohere in the right way with one's other beliefs (and other doxastic attitudes). So the requirement to believe what one's evidence supports and the various requirements to be coherent ultimately come to the same thing.

A different kind of implicit answer (which I judge to be more popular than the first, at least recently) takes it that so-called "requirements" of coherence are ultimately reducible to facts about what the evidence supports. The idea here is that it is a feature of the correct theory of evidential support that one's evidence can never support attitudes that are jointly incoherent.[5] This is easiest to illustrate with respect to the requirement of non-contradiction. Whatever some proposition p is, and whatever one's evidence is, it seems hard to see how one's (total) evidence could simultaneously support believing p, and also support believing not-p. So, if rationality is a matter of taking the attitudes that one's evidence supports, it will just fall out as a consequence (so the argument goes) that it is never rationally permissible to be in a state such that one has attitudes that violate certain coherence constraints. On this view, there are no *sui generis* coherence requirements, or (to be slightly more metaphysically non-committal) if there are, they are normatively superfluous, since they do not require anything that evidence-responsiveness didn't require already. Call this view 'Superfluousness'. Given Superfluousness, for every would-be coherence requirement of the form *don't simultaneously have attitude A and have attitude B*, there is a corresponding metaphysical constraint on what combinations of facts about evidential support are possible, of the form *one's evidence cannot decisively support both attitude A and attitude B*, and it won't do any harm to move between these two claims freely.

The two views just sketched clearly differ in a number of crucial respects, but both share the feature of ultimately trying to deny the thought that there

---

[2] C.f. Kolodny (2005, 2007); Fogal (ms.); Lord (2014); Sylvan (ms.); and Easwaran & Fitelson (2015).

[3] C.f., e.g., BonJour (1985).

[4] C.f., e.g., Titelbaum (2010).

[5] Something like this strategy is appealed to by Kolodny (2005) and Lord (2014). It also seems to be at work in Huemer (2011: 7; compare also p. 3 with p. 4, n. 7).

are really two different, important normative phenomena here to be pulled apart. On the first view, evidence-responsiveness and coherence ultimately come to the same thing, so there simply is nothing to pull apart. On the second view, evidence-responsiveness entails coherence, but not vice-versa. At the *most*, then, we could simply distinguish a "more demanding" sense of the term 'rationality'—on which it requires evidence-responsiveness *and* (as an automatic consequence) coherence—from a "less demanding" sense of the term—on which it requires coherence only. But at least to some philosophers, it's not apparent what the interest of the second notion would be. Even if there is an important intuition that rationality requires coherence between attitudes, the first notion of rationality can capture that intuition, whilst doing more besides. There is, on this view, no obvious reason why we need to make room for a notion of rationality as coherence alone.

The aim of this paper is to make the case that this sanguinity is misplaced. This is so because, as I will argue, coherence and evidence-responsiveness can in fact come into *conflict*.[6] That means that a reduction in either direction will not work, and neither coherence nor evidence-responsiveness will guarantee the other; Superfluousness is false. As I will argue, we must instead acknowledge coherence requirements and facts about what the evidence supports as distinct normative phenomena.[7]

More specifically, I will argue that in cases of all-things-considered misleading higher-order evidence—cases where one's evidence all-things-considered supports some attitude, but also all-things-considered supports believing that it does not support that attitude—there is a conflict between believing what one's evidence supports and satisfying a requirement that I call "inter-level coherence". It is plausible both that there can be such cases of all-things-considered misleading higher-order evidence, and that the requirement of inter-level coherence is a genuine coherence requirement of rationality (that is, not only is it of the right category to be a *coherence* requirement, but it is also a genuine *requirement of rationality* in the sense that in violating it, one is thereby irrational). A number of philosophers have recently noticed a puzzle rather like the one that I will consider, and inferred that we must either deny the requirement of

[6] This is a revisionary view even relative to those few who do distinguish coherence requirements from evidence-responsiveness, and acknowledge the need to theorize the former. So Easwaran & Fitelson (2015), who nicely distinguish substantive "norms" of evidence-responsiveness from coherence "requirements", assume that it is a decisive objection to a coherence requirement if satisfying it would require one to violate those substantive norms—thus taking it for granted that conflicts between the two are impossible.

[7] I'm using the term "normative" broadly here. On another usage of "normative", a requirement is only normative if there is necessarily *reason* to comply with it. (C.f. the debate between Kolodny (2005), Raz (2011: ch. 8), Broome (2013: ch. 11), and others.) I leave it open whether coherence requirements are normative in this sense.

inter-level coherence or deny that one can have all-things-considered misleading higher-order evidence. As I will argue, that is a false dichotomy. We need deny neither claim if we distinguish coherence requirements from claims about what the evidence supports. So, on the way to my conclusion, we will also discover how to dissolve this apparent puzzle about higher-order evidence.

It's an upshot of my argument that we need distinct normative concepts to talk about coherence and evidence-responsiveness. On one possible way of talking—highly revisionary relative to epistemology but far less so relative to the practical rationality literature—we could restrict the term 'rationality' to the satisfaction of coherence requirements alone, using the terminology of *reasons* to talk about evidence and evidence-responsiveness.[8] I will not insist on this terminology here. But one way or another, the distinction between coherence and evidence-responsiveness must be marked. At the moment, the usage of 'rationality' in epistemology is frequently indeterminate between referring to evidence-responsiveness (i.e., epistemic reasons-responsiveness), referring to conformity to coherence requirements, or referring to both. If evidence-responsiveness and coherence can never come into conflict, this indeterminacy may not matter much. But if they can, the indeterminacy will lead us into serious, and substantive, philosophical mistakes.

I. A puzzle to motivate the argument

Consider the following three claims. Throughout, D(p) is a possible doxastic attitude (D) of a subject (S) towards a proposition (p): believing p, disbelieving p, or suspending judgment about p.[9]

---

[8] For those that sharply separate reasons and rationality in the practical context, see, e.g., Broome (2013: chs. 5–6); Scanlon (1998: ch. 1, 2007); Setiya (2004); Kolodny (2005); McDowell (1998: ch. 5); Ridge (2014: ch. 8); Rawls (1996: 50–51); Smith (2007). Interestingly, those that resist this sharp separation in the practical domain still tend to recoil from equating rationality with responding to reasons *simpliciter*. Rather, they say that rationality consists in responding to *beliefs* about reasons, or that it consists in responding to some privileged *subset* of one's reasons. See, e.g., Schroeder (2009); Raz (2011: ch. 5); Parfit (2011: ch. 5); Lord (2014). It's interesting to ask whether the asymmetry between the epistemic and practical literatures here is simply a sociological quirk of different word-usage, or whether it reflects some kind of deeper substantive difference in orientations. Elsewhere (Worsnip forthcoming), I suggest that the latter diagnosis is correct, and argue that such a deeper asymmetry cannot be vindicated.

[9] So our focus here is on outright attitudes, and not on graded attitudes (i.e., credences). There are analogues of both (ER) and (ILC) for credences. But formulating the credence-specific version of (ILC) turns out to be very complex and to raise a host of its own issues. So to keep things manageable, I focus on outright attitudes here.

Evidence requirement (ER). If S's evidence supports D(p), then rationality requires of S that she takes D(p)

Inter-level coherence (ILC). Rationality requires of S that
(i) S believes that her evidence supports D(p) → she takes D(p)
(ii) S believes that her evidence does not support D(p) → she does not take D(p)

Possibility of iterative failure (PIF). It is possible that:
(i) S's evidence supports D(p); and
(ii) S's evidence supports believing that her evidence does not support D(p)

(ER) is the claim that rationality requires one to take the attitudes that one's evidence supports, spelled out. (ILC) is a coherence requirement governing the relations between one's first-order beliefs and one's higher-order beliefs. (PIF) is the claim that one can have all-things-considered misleading higher-order evidence, such that one's evidence supports a false claims about what one's evidence supports (I am calling such instances cases of "iterative failure", since they are cases where one's evidence fails to iterative across levels of beliefs).

Suppose now that (PIF) obtains with respect to some doxastic attitude D towards some proposition q: that is, S's evidence supports D(q), but S's evidence supports believing that her evidence does not support D(q). Then, S cannot satisfy both (ER) and (ILC). To see that, consider the following table:

	S takes D(q)	**S doesn't take D(q)**
S believes that her evidence does not support D(q)	S violates (ILC)	S violates (ER)
S does not believe that her evidence does not support D(q)	S violates (ER)	S violates (ER) (twice over)

Either S takes D(q), or she doesn't. And either S believes that her evidence does not support D(q), or she doesn't. So S must be in one of the four boxes. Yet in each of the four boxes, S violates either (ER) or (ILC). So if there are cases of iterative failure—as (PIF) says there are—then by the same token, there are cases where (ER) and (ILC) are not co-satisfiable.

Something like this puzzle—not quite in the form that I have spelled it out here—has been noticed before. As far as I can see, the existing responses in the literature fall into two broad categories:

1. Deny that cases of iterative failure are possible (that is, deny (PIF)). So conflicts between (ER) and (ILC) cannot arise, at least not in this way.[10]

2. Deny that (ILC) is a genuine requirement of rationality, and hold that in cases of iterative failure, one can violate (ILC) while remaining rationally flawless.[11]

I am going to argue that both of these responses fail. (ILC) and (PIF) are both independently plausible, I will argue, and neither provides principled reason for denying the other. Most philosophers who discuss our puzzle seem to assume that options 1 and 2 are the only options. If this assumption were true, such that we simply must deny either (ILC) or (PIF), then the truth of one would of course then provide very strong grounds indeed for denying the other. However, the assumption is too hasty. There are at least two further options available to us. The first (option 3 overall), not unequivocally endorsed by anyone in response to this particular puzzle, but in the spirit of some remarks made by at least one philosopher,[12] is:

3. Maintain that (ER) and (ILC) are both requirements of rationality (in the same sense of 'rationality'), and hold that cases of iterative failure are rational dilemmas, whereby, whatever one does, one violates a requirement of rationality.

The last option, and the one I favor, is:

4. Maintain that (ER) and (ILC) are, properly understood, fundamentally different kinds of normative claim, such that they should not be stated using the same normative concept.[13]

As I said in the introduction, one way to execute this strategy is to say that only (ILC), and other coherence requirements like it, are requirements of

[10] Feldman (2005); White (2007: 120); Huemer (2011); Greco (2014); Titelbaum (2015). Horowitz (2014) takes a nuanced stance on (ILC)—arguing that it holds for the most part, but occasionally fails. Nevertheless, like the other authors mentioned here, she assumes that any admission of (PIF) requires admitting a corresponding failure of (ILC).

[11] Coates (2012); Lasonen-Aarnio (ms.); Weatherson (ms.); and arguably Wedgwood (2012).

[12] Christensen (2007, 2010, 2013).

[13] This option does not, as far as I can see, even feature in existing taxonomies of the options in response to our puzzles and related ones. Compare, e.g., Huemer (2011: 5); Greco (2014: 206); Titelbaum (2015: 278–279).

rationality, strictly speaking. As such, we should reject (ER) as it is currently stated, and replace it with:

> **(ER*)** If S's evidence supports D(p), then S has most epistemic reason to take D(p)

However, as I already said in the introduction, there are different ways to use the relevant terminology here. The most important feature of strategy (4) is that it affirms that (ER) and (ILC) are different kinds of normative claim.

I want to stress at the outset that there are a number of commonalities between option 3 and option 4. Perhaps most importantly, both strategies accept that (genuine) coherence requirements can come into conflict with evidence-responsiveness. The difference between the two strategies is that while option 3 construes this as a conflict between two different requirements of the fundamentally same kind (two different requirements of rationality), option 4 construes it as a conflict between two fundamentally different kinds of normative demand (on my preferred terminology, that of satisfying the requirements of rationality, and responding to one's reasons). Later, in section VI, I will say something more about this difference, and offer some reasons to prefer option 4 over strategy 3. However, on *either* view, it's crucial to see that the following popular claim turns out to be false:

> **Evidentialism.** S's doxastic states are rational iff they satisfy (ER).

Even on option 3—which allows that (ER) and (ILC) can be requirements of rationality in the same sense of the term 'rationality'—Evidentialism is false. This is because, if satisfying (ER) does not suffice for satisfying (ILC), and (ILC) is still a genuine requirement of rationality, then satisfying (ER) does not suffice for rationality. This is already a highly significant result, even if one does not go as far as option 4.

The plan of attack for what remains is as follows. In part II, I will issue some further preliminary clarifications about some of the relevant notions in play. In part III, I will give an explication and partial defense of (ILC). In part IV, I will give an explication and defense of (PIF). Having argued that cases of iterative failure are possible, I will in part V complete my defense of (ILC) by arguing that its plausibility is not diminished in cases of iterative failure specifically. That will complete my case that (genuine) coherence requirements can come into conflict with evidence-responsiveness. Parts VI turns to the issues that remain once we have already accepted that, arguing for option 4 over option 3—that is, in favour of distinguishing the kinds of normative demands in play, rather than describing the situation as

one where demands of the *same* kind come into conflict. Part VII concludes.

II. Further terminological clarifications

In this section, I will quickly say a little more about the notions of 'coherence' and 'evidential support' that I am working with, and then finally about the term 'justification', which I will not use much.

On my usage of the term 'coherence', coherence requirements are distinctively about which attitudes (or absences of attitudes) it is rationally forbidden (or permitted) to *combine* together. In general, coherence requirements do not say that individual attitudes are mandated or forbidden. So, in a case where one violates a coherence requirement—holding two attitudes that are jointly irrational—the coherence requirement itself does not say which attitude should be revised. Rather, it just says that the conflict must be eliminated somehow. Other normative considerations external to the coherence requirement itself may or may not settle how such a resolution should proceed.

This usage of 'coherence' is relatively broad. Two attitudes need not be deductively inconsistent to be incoherent in my sense. Nor is 'coherence' restricted to the *probabilistic* "coherence" of credences. However, when I talk about incoherence, I mean *genuine* incoherence. Some philosophers may use 'incoherent' to stipulatively refer to attitudes that violate certain putative coherence requirements: for example, to credences that do not conform to the probability axioms, or to beliefs that are not consistent or closed under logical implication.[14] But it is, in my sense of 'coherent', a substantive question for debate whether such attitudes *are* incoherent or not. Someone who wants to hold that these combinations of attitudes are not irrational may hold that they are not properly speaking incoherent, rather than endorsing the idea that (genuine) incoherence can be rational.

Next, the notion of evidential support. What the evidence supports can be understood as the attitude or attitudes that you have most (epistemic) reason to take, given your total evidence. Given this gloss on evidential support, it should be clear that we will do better to talk of evidential support for *attitudes towards propositions* than to talk (as some do) of evidential support for *propositions themselves*. One has a reason to take an attitude towards a proposition, not a reason for a proposition. This way of talking is better for another reason. When it comes to propositions themselves, there are only two candidates for being supported: p and its negation. But when it comes to attitudes towards some proposition p, there are *three* (outright) attitudes that can be supported by one's

[14] C.f., e.g., Foley (1993: §4.6) and possibly Pryor (ms.).

total evidence: believing p, disbelieving p (believing not-p), and suspending judgment about p. It is important that we don't lose sight of the possibility of one's total evidence supporting suspending judgment. In a case where the evidence in favor of (believing p) is only slightly weightier than the evidence against (believing) p, we do not want to say that one's total evidence supports believing p; on that interpretation, (ER) is obviously false. Rather, in such a case, one's total evidence supports suspending judgment. This fits with my gloss on evidential support, since given one's total evidence, one has most reason to suspend judgment on p. Talking of support for attitudes (rather than propositions) makes clearer room for the full range of possible doxastic responses, and for support for each of them.[15]

'Support' can also be used in either a *pro tanto* or *pro toto* sense—we can refer to evidence as providing partial, defeasible *pro tanto* support for a particular doxastic attitude, or to the doxastic attitude that your evidence supports *pro toto*—that is, on balance. As my gloss above makes clear, I am using the term 'support' in the *pro toto* sense (unless I specifically indicate otherwise). The attitude supported by the evidence is just the attitude that one has *most* reason to take, given one's total evidence. If the evidence *pro tanto* supports the different possible attitudes to various degrees, there must be one (or more) attitudes that it supports the most. So, given my usage, there will always be at least one attitude that the evidence supports with respect to a proposition. If your evidence does not support believing p or support disbelieving p, then it supports suspending judgment about p.[16] So, for example, a case where you have very scant evidence both for and against some proposition is not a case where your total evidence doesn't support *any* doxastic attitude; rather, it is one where your total evidence supports suspending judgment.

---

[15] An anonymous referee suggested that individual items of evidence never themselves provide reason to suspend judgment (as argued by Schroeder (2012)). I cannot settle that here, but even if it is right, I think it still makes sense to talk of a *total* body of evidence, through its structural features (for example, its consisting of quite a lot of evidence in favor of a potential belief, and quite a lot of evidence against it), supporting suspension.

[16] Given this usage, (ER) entails the claim that if S's evidence does not support D(p), then rationality requires of S that she does not take D(p). The entailment runs in both directions for cases in which a subject must either believe, disbelieve, or suspend judgment, so if there are some such cases, my argument will also extend to this more negative analogue of (ER).

What about the possibility that your evidence supports more than one different doxastic attitude equally well?[17] My usage of 'support' does not exclude this possibility, since the reasons for taking two or more different attitudes might be equal. Anyone who is attracted to this possibility, however, will not like (ER) as it is currently formulated, since we definitely don't want to say that rationality requires you to take *each* of the doxastic attitudes that your evidence supports equally well. Rather, those who think this situation is possible should read D (as it occurs in every principle that I refer to) as the disjunction of the different attitudes that the evidence supports. My arguments will be unaffected by this modification.

Last, 'justification'. In my view, the term 'justified belief', in contrast to 'rational belief', is best used simply to refer to a belief's being supported by the evidence.[18] Some philosophers may (intelligibly) use 'justification' instead to refer to the property or properties that are—in addition to truth, and perhaps anti-Gettier luck—required for a belief to count as knowledge.[19] Whether that property is equivalent to the property of being supported by one's evidence is a substantive question that I will not take a stand on here. My topic is the relationship between the property of coherence and that of being supported by the evidence. The relationship of these properties to the knowledge-making property is an independent question.

---

17 This may put readers in mind of the debate between 'uniqueness' and 'permissivism': see, e.g., White (2005); Ballantyne & Coffman (2011). However, these authors tend to treat the thesis that there is always one unique attitude that the evidence supports interchangeably with the thesis that there is always one unique attitude that is *rational* given a subject's evidence. These two formulations are only equivalent if we assume that rationality is a matter of evidence-responsiveness. I reject uniqueness as a thesis about rationality, but am agnostic on it as a thesis about what the evidence supports.

18 Many epistemologists have traditionally used 'justified' and 'rational' interchangeably, and perhaps also assumed that both are equivalent to the knowledge-making property. But recent theorists have developed various intelligible distinctions between rationality and justification; see, e.g., Goldman (1986, 2009), Pryor (2004: 363–365), Littlejohn (2012), Sylvan (ms.). My main point here is that there are various different underlying properties to be marked: the property of being coherent, the property of being evidence-responsive, the knowledge-making property. However one uses terms, the thesis that any of these properties are co-extensive with one another will be a substantive one open to debate.

19 This may be how Goldman (1986) uses the term when he says he claims that justification is not a matter of being supported by the evidence but of "external" reliability. Note that this contrast implicitly depends upon assuming that evidential support cannot itself be construed in externalist terms, an assumption that has been increasingly rejected more recently (e.g. Williamson (2000); Pritchard (2002)). Note also that, since (as noted in the previous footnote) Goldman explicitly distinguishes justification and rationality, he may still follow the vast majority of epistemologists in identifying *rationality* with evidence-responsiveness.

III. (ILC) explicated and (partially) defended

Here, once again, is the inter-level coherence requirement:

> **Inter-level coherence (ILC).** Rationality requires of S that
> (i) S believes that her evidence supports D(p) → she takes D(p)
> (ii) S believes that her evidence does not support D(p) → she does not take D(p)

(ILC) forbids certain combinations of first-order and higher-order attitudes. 'Rationality requires' takes "wide scope" over the conditionals in (i) and (ii), which (as a matter of stipulation) are material.[20] So (ILC) bans, for example, simultaneously believing p, and believing that one's evidence does not support believing p. This wide-scope form is important. It may be tempting to think of (ILC) as a kind of 'subjectivization' of (ER), replacing the notion that rationality is about conforming your beliefs to the evidence with the idea that rationality is about conforming your beliefs to what you believe to be the evidence.[21] But in an important way this is misleading. (ILC) is not the following:

> **Narrow scope inter-level coherence (NILC).**
> (i) S believes that her evidence supports D(p) → rationality requires of S that she takes D(p)
> (ii) S believes that her evidence does not support D(p) → rationality requires of S that S does not take D(p)

As least read naively, this narrow scope requirement is implausible. It entails that, if you have crazy second-order beliefs about what your evidence supports, then rationality *requires* you to have correspondingly crazy first-order beliefs. Moreover, (NILC) is just obviously incompatible with (ER). For in any situation where someone has mistaken beliefs about what their evidence supports, (ER) and (NILC) will issue contradictory requirements. A fan of (ER) will feel happy and principled in rejecting (NILC).

By contrast, (ILC) does not have these same consequences. Because (ILC) is wide-scope, it is neutral on *how* to resolve inter-level mismatches. Consequently, it does not give any special authority to one's beliefs about one's evidence. They are just as open to revision as one's first-order beliefs. Suppose you violate (ILC), so that there is a mismatch between your first-order attitudes and your higher-order beliefs about what your evidence supports. You can come to satisfy (ILC) *either* by revising your first-order

---

[20] For the idea that coherence requirements in general are "wide-scope", see Broome (1999, 2013: ch. 8); Brunero (2010); Dancy (2000: 60–76); Scanlon (2007); Wallace (2001); Way (2010).

[21] C.f. the criticism of (ILC) in Wedgwood (2012: 291); see also White (2007).

attitude, or by revising your higher-order belief. (ILC) does not require either revision in particular, and leaves it open that in some particular case you ought (for other reasons) to satisfy it by one kind of revision rather than the other kind. All it says is that you are required to revise *something* so as to eliminate the conflict in your doxastic states. Consequently, one might think initially that (ER) and (ILC) will always be co-satisfiable. The fan of (ER) can hold that when one has mistaken beliefs about what one's evidence supports, one should come to satisfy both (ER) and (ILC) by revising one's beliefs about what the evidence supports.[22] As we will see, it is only when we bring in (PIF) that the two putative requirements come into conflict.[23]

These clarifications help us to see what the motivation for (ILC) is not. Cases of iterative failure make it clear that one can be blamelessly misled about what one's evidence supports. One might think that (ILC) is motivated by an attempt to seek a luminous epistemic property—one to which one has a kind of infallible epistemic access, or cannot be blamelessly misled about. This can be thought of as an extreme form of epistemic 'internalism'. And the thought might be that although one can be blamelessly misled about what one's evidence supports, one cannot be blamelessly misled about what one believes one's evidence supports. But this is not my thinking: in fact, I reject these claims. It is part of the spirit behind (PIF) that one can be misled about just about anything interesting, including one's own beliefs.[24] In my view, to say that one can be misled about one's evidence, but not about one's beliefs about one's evidence, represents a kind of unstable midway position.

(ILC)—or variants of it close enough to preserve the validity of my argument—is widely accepted.[25] So if (ILC) does not rely on any extreme

---

22 See Feldman (1988: 412).

23 Note also that (ILC) stays silent about situations in which you suspend judgment about what your evidence supports. Horowitz (2014: 724) describes this view as allowing for the rationality of "moderate mismatches" between one's first-order and higher-order beliefs (see also Feldman (2005: 118 n. 6); Huemer (2011: 1); Hazlett (2012: 218–219)). I think this is misleading. Remember that (ILC) *does* require that if one believes that one's evidence supports suspending judgment about p, then one suspends judgment about p. The case where it stays silent is that where one suspends judgment about what doxastic attitude toward p one's evidence supports, *including* about whether one's evidence supports suspending judgment about p. It's not like your higher-order attitude commits you to suspending judgment on the first-order question here. So I don't think there's any mismatch here at all, even a "moderate" one.

24 See, amongst others, Williamson (2000: esp. ch. 4); Schwitzgebel (2008); Srinivasan (2013). For an application of these points to draw lessons about norms epistemic and otherwise, see Hawthorne & Srinivasan (2013); Srinivasan (forthcoming).

25 Harman (1986: 39); Bergmann (2005); Elga (2005); Feldman (2005); Kolodny (2005: 521); Scanlon (2007); Christensen (2007, 2013); Smithies (2012: 283–284) Broome (2013: 98); Greco (2014); Horowitz (2014); Titelbaum (2015).

internalism, what does motivate it? Let's begin with Adam Elga's (2005: 115) presentation of a case where (ILC) is violated, and his reaction:

> My friend Daria believed in astrology. For example, she thought that because of her astrological sign she was going to be particularly lucky over the next few weeks. That was bad enough. But when I tried to persuade her that astrology is unfounded, I discovered something even worse. I gave Daria evidence against astrology—studies showing that the position of the distant stars at the time of one's birth has no bearing on one's personality or prospects. Daria agreed that the studies were significant evidence against the truth of astrology, and that she had no countervailing evidence of comparable strength. But that was not the end of the matter. "I still believe in astrology just as much as I did before seeing the studies," she said. "Believing in astrology makes me happy."
>
> I was floored. Daria's original belief in astrology was less than perfectly reasonable. But this—believing in astrology even though by her own lights the evidence went against it—was an insult to rationality.

Notice that Elga here takes (ILC) to be a more fundamental and non-negotiable part of rationality than (ER). In calling Daria's original belief "less than perfectly reasonable," he stops short of calling it *irrational*—it is her violation of (ILC) that is, for Elga, "an insult to rationality". The idea here is that Daria is criticizable for two reasons in this example. She is criticizable for having a belief that does not accord with her evidence. But she is also criticizable—and perhaps in a different way—for having this belief in the face of her own judgment about what her evidence supports.

Why should this second state of Daria's be criticizable? One way to motivate the thought is to appeal to an analogy with the so-called 'enkratic requirement' of practical rationality. This requirement forbids instances of akrasia, where you fail to intend to do what you believe you ought to do.[26] Several writers who defend something close to (ILC) have presented it as a requirement that forbids 'epistemic akrasia'.[27]

At least as I have formulated it, (ILC) is not quite an anti-akrasia requirement, strictly speaking, because it refers to what one believes one's evidence supports, and not to what one believes one ought to believe. This is not an accident: in a moment, I will give an argument for (ILC) that works only for a formulation in terms of evidential support.[28] Still, at least if the agent implicitly takes what her evidence supports to be determinative of

---

[26] See Broome (2013).

[27] See Horowitz (2014); Greco (2014); Titelbaum (2015).

[28] And see Broome (2013: 88–89, 92–96) for some considerations *against* a formulation in terms of what one ought to believe.

what she ought or has most reason to believe, there is a sense in which a violation of (ILC) amounts to a kind of akrasia. The idea here is that it is irrational to fail to comply with one's own judgments about what one has reason to believe; it is a failure to match up to one's own standards; in some sense a failure "by one's own lights".[29]

Note that this way of motivating (ILC) actually, far from denying the normative importance of what one's evidence actually supports (as a simple 'subjectivization' of (ER) might be thought to), actually *trades* on this normative importance. As such, (ILC) should be attractive to those who appreciate the normative importance of evidence: violating (ILC) amounts to what is, from the subject's own point of view, an act of disrespect for her evidence. That is one reason why (ILC) has actually been accepted by epistemologists that take evidence to be normatively fundamental.[30]

At the same time, these observations already start to point us to some of the salient differences between reasons and coherence requirements. What one has most (epistemic) reason to believe is just constituted by what one's actual evidence supports, and the facts about what that evidence actually supports. But what involves the agent in *incoherence* is that she *takes herself* to be failing to believe what her evidence supports. These facts are, at least arguably, related: if (actual) evidence has a constitutive tie to epistemic reasons, then it may consequently be incoherent to *take oneself* to be believing against it. But they are nevertheless distinct. The idea is not that what one *takes* to be one's evidence is determinative of what one has *reason* to believe, any more than the enkratic requirement claims that what you believe you ought to do is determinative of what you ought to do. Nor is the idea that there is really some kind of internal incoherence in failing to

---

[29] However, the language of failure "by one's own lights" is slippery and thus arguably problematic. (This is brought out both by Broome (2013: 91–93), a proponent of the enkratic requirement, and by Lasonen-Aarnio (ms.: 11), a critic of it.) The idea that incoherence must amount to failure by one's own lights may seem to suggest the idea that to be incoherent, one must *accept* that one is being incoherent. But this is liable to lead to a vicious regress, where no coherence requirement can convict as irrational those who reject it as a genuine requirement of rationality, and so no coherence requirement can be true in full generality. Proponents of coherence requirements should hold, instead, that any violation of a genuine coherence requirement is incoherent, regardless of whether the agent accepts or rejects any theoretical proposition about the coherence requirement itself. This may be compatible with holding that violations of coherence requirements are, in some sense, "failures by one's own lights"—but this cannot mean that one has to take oneself to be incoherent in order to be such. Moreover, as an anonymous referee pointed out to me, in some contrastingly broad sense of "failures by one's own lights", even failures of evidence-responsiveness will be "failures by one's own lights", since one's evidence is on any account constrained by one's perspective.

[30] See esp. Feldman (2005) and Adler (2002). Many if not all of the other theorists cited as supporting (ILC) in fn. 25 are in this category.

believe what one's evidence *actually* supports. Rather, reasons and coherence requirements have to be pulled apart here. Much more on that later.

Some may be unmoved by the appeal to analogy with the enkratic requirement. They may reject that requirement. Fortunately, there is another, independent way to motivate (ILC) that does not appeal to the analogy with the enkratic requirement. In fact, it may put the former on a stronger footing than the latter, and to see this we can consider a *contrast* with the practical case. In the case of ordinary practical akrasia, an agent believes she ought to Φ, but does not intend to Φ. While this may seem irrational, being in such a state is easily intelligible, even first-personally: it's all too easy to see how one could believe one ought to Φ, but not be able to bring oneself to have the corresponding intention. On the other hand, the state of believing that one's evidence does not support believing p, but nevertheless believing p, is harder to make sense of, at least from a first-personal perspective. It amounts to saying "I have nothing that gives any adequate indication to me that p is the case; nevertheless, p *is* the case".[31] That, it seems to me, is a deeper kind of incoherence than practical akrasia. It is not impossible for an agent to be in such a state, but it is hard to fully make sense of. First-personally, these states do not seem capable of withstanding serious reflection. And third-personally, while we can imagine such agents, in describing and explaining them we reach for some story involving self-deception or a failure to recognize their own mental states.

I think an explanation of these facts can be given. To believe something is to take it to be true. But evidential reasons for belief are precisely those reasons that bear on how likely the belief is to be true. As such, to believe while judging oneself to lack sufficient evidence amounts to holding that p is true, but also isn't especially likely—in light of *all* the available information[32]—to be true. That is a deep kind of incoherence; one associated with beliefs that defy one's judgments about *evidential support* specifically, and not just associated with the general phenomenon of akrasia. Whereas I can, in a fully clear-eyed way, have reflective thoughts like "I shouldn't watch the game, but I will", it's less clear how, in a fully cleared-eyed way, I can have reflective thoughts like "the game starts at 3pm, but is (in light of all the information I have) unlikely to start at 3pm".

---

[31] As others have noted, this sort of claim is Moore-paradoxical. See also Adler (2002), Feldman (2005), Huemer (2011) and Smithies (2012).

[32] This qualification is important, since there's nothing incoherent about believing something that one takes to *have been* unlikely given some restricted background set of information.

If this starts to make (ILC) sound like it plays a descriptive (as contrasted with normative) role in understanding our psychology: I think that claim is partly right. But the claim is not that (ILC) is impossible to violate.[33] It is, at most, that it is hard to violate without some kind of failure of transparency to oneself—and that to do so is to make oneself hard to interpret or make sense of.[34] These facts may be related: if one fully recognizes that one violates (ILC), one will find oneself hard to make sense of: that is perhaps part of why it is hard to come to this recognition and to sustain the offending mental states in the face of it. Again, this foreshadows some important contrasts of kind between coherence requirements and reasons. (ILC)—and perhaps many other coherence requirements too—plays a role in the interpretation and attribution of mental states that reasons do not: to interpret others, we will tend to assume that they at least generally, and at least under conditions where they are reflectively aware of their own mental states, will tend to satisfy (ILC).

That does not mean that (ILC) does not also play a normative role, but its normative role is somewhat different from that of reasons. To get a grip on the distinctive irrationality of violating (ILC), one needs to get into a first-personal way of thinking and see how hard it is to reflectively sustain the states that violate it. If we *begin* with the thought that third-personally, we can make sense of situations in which one's evidence supports having states that violate (ILC), and then directly conclude that one can violate (ILC) without any irrationality whatsoever, we overlook the distinctive first-personal irrationality of incoherence. Such an approach just takes the view that I earlier called Superfluousness for granted—taking it that any coherence requirements that there are will just fall out of facts about evidential support, and have no distinctive status of their own. But I suggest that this is a mistake. Even if, since one is in a case of iterative failure, violating (ILC) does not *actually* involve believing against one's evidence, the violation will still involve *taking* oneself to believing against one's evidence, and that will be incoherent for all the same reasons, irrespective of what one's evidence actually supports.

For these reasons, I don't think that simply producing examples of iterative failure—as opponents of (ILC) have done[35]—suffices to put pressure on the positive rationale for (ILC). We'll only be able to fully secure this point, however, when we have cases of iterative failure on the table. So I will return to it in section V. Before that, it is time to defend (PIF).

---

[33] Some have been tempted by that stronger claim; see especially Adler (2002) and Hurley (1989: 130–135, 159–170).

[34] I hope to develop this idea in future work.

[35] See, e.g., Coates (2012); Lasonen-Aarnio (ms.); Weatherson (ms.).

IV. (PIF) defended

By now, you may well be wondering whether denying (PIF) might be the right way to go. Let us remind ourselves of (PIF):

> **Possibility of iterative failure (PIF).** It is possible that:
> (i) S's evidence supports D(p); and
> (ii) S's evidence supports believing that her evidence does not support D(p)

Denying (PIF) requires denying that one can have all-things considered misleading evidence about what one's evidence supports; in other words, justified false beliefs about what one's evidence supports are impossible. That is a strong claim, as any claim that a particular kind of justified false belief is impossible would be.

Let me be clear that in both parts (i) and (ii) of (PIF) above, it is S's *total* evidence that is being referred to. The claim is not just that S's first-order evidence *alone* can support D(p), while S's higher-order evidence *alone* supports believing that her evidence does not support D(p). That is a fairly obvious claim, and would require little argument. But it would not be enough for our purposes. For (ER) requires a subject to take the attitudes that her evidence as a whole *pro toto* supports. So, throughout, read (PIF) as pertaining to a subject's *total* evidence. That's what I mean by talking of "all-things-considered" misleading evidence. That said, we can still distinguish between what a subject's total evidence supports *at the level of her first order attitude* (D(p)), and what her total evidence supports *at the level of her higher-order attitude* (her belief as to whether her evidence supports D(p)). And it may be—as I'll suggest—that her first-order evidence has *greater weight* (relative to the rest of her total evidence) with respect to this first question, while her higher-order evidence has greater weight with respect to the second question.

I assume that we can think of one's evidence as consisting in some body of items, whether those be propositions, or facts, or even (if you prefer) sensations and appearances. One's total body of evidence then supports certain doxastic attitudes. Now, it's important to distinguish two quite different ways to be misled about what one's evidence supports. On one hand, there is the possibility that one might be misled about which items are part of one's body of evidence: that one might be led to think that some proposition (or fact, or appearance) is part of one's evidence when it actually isn't, or led to think that some proposition (or fact, or appearance) isn't part of one's evidence when it actually is. In this case, one is misled about what one's evidence *is*. On the other hand, there is the possibility that one is misled not about what one's evidence is, but about the *evidential support relations* that hold between particular items in one's body of evidence and potential doxastic attitudes.

One can be misled about these support relations even if one identifies which items are *part* of one's evidence entirely correctly; similarly, one can be mistaken about which items are part of one's evidence even if one is omniscient about which (potential) items support which attitudes, and how strongly. Both of these ways of being misled will eventuate in being misled about what one's total body of evidence supports, and so to the possibility of iterative failure. But the source of the mistake will be different in the two cases. In one case, the source of the mistake is one's misleading evidence about what one's evidence *is*. In the other, the source of the mistake is one's misleading evidence about the evidential support relations between one's evidence and potential doxastic attitudes.

In a moment, I'll take these cases one at a time, and try to explain how iterative failures of either kind can come about. Before that, a point about method. I will try throughout this section to give intuitive cases in which it is plausible that iterative failure occurs. However, there is a very general reason why it is difficult to give knockdown intuitive cases of such iterative failure, namely that it is always a somewhat vague and hard-to-identify matter just how much evidence is required to decisively support a doxastic attitude. Cases of iterative failure exploit the idea that sometimes your evidential support for a first-order attitude is stronger than the support for some corresponding higher-order attitude, or vice versa. Thus, in cases at the margins, it may be enough to *decisively* support the first-order attitude, without being enough to tip the balance of evidence at the level of higher-order attitudes. But because of the difficulty of determining how much evidence is decisive, it's always possible in any concrete case to protest that it strikes one that the evidence is not enough to decisively support the first-order attitude after all, or (perhaps) that it is enough to support the corresponding higher-order attitude. (This same problem arises equally for the more familiar project of giving cases of the failure of the "KK" principle, which I also discuss below.) For these reasons, at times I will need to supplement my cases with theoretical arguments to the effect that we should not always expect things to evidentially line up across levels in the way that the opponent of iterative failure wants. The concrete cases thus serve to illustrate in what *kinds* of cases iterative failure may occur, even if one can quibble over the particular cases given.

(a) Cases of being misled about what one's evidence is

I'll begin by illustrating a method for generating cases of iterative failure out of failures of the so-called "KK principle", on which if one knows, then one knows that one knows.[36] It shouldn't be surprising that these two

[36] See Williamson (2000: esp. ch. 5) for classic arguments against KK.

possibilities are linked. If KK fails, knowledge fails to iterate across levels; if that can happen with knowledge, it's not surprising that it can also happen with evidence and evidential support.

Let's begin with a classic case of KK failure, that of the unconfident examinee.[37] In this case, a high-school student, Sayeqa, is taking a history exam, and is asked for the dates of various English monarchs' ascendance to the throne. For each monarch, Sayeqa gives the correct answer. Her doing so is based on her having learnt the dates in a normal way, through instruction in a classroom, and coming to associate the correct dates with each monarch. What she is in fact doing is recalling these dates based on this association. Yet she is of a nervous sort of disposition in test environments, and this disposition causes her to feel as if her answers are guesses.

Take the proposition that Charles I ascended to the throne in 1625—one correct answer that Sayeqa has given in a long list of many. The verdict that KK deniers take as natural in this case is that Sayeqa knows that Charles I ascended to the throne in 1625, but that she does not know that she knows this. However, it would be wrong to say that Sayeqa's knowledge that Charles I ascended to the throne in 1625 somehow amounts to her knowing without possessing evidence that supports believing this proposition. She does possess evidence: the evidence she gained in the classroom.[38] That evidence decisively supports believing that Charles I ascended to the throne in 1625. What Sayeqa lacks is decisive evidence *that* she has this evidence—just as she lacks knowledge that she knows. But in the absence of decisive evidence that she has this evidence, and in the presence of the (misleading) impression that her answer is just a guess, it seems that Sayeqa's evidence supports believing that her evidence supports suspending judgment about the proposition that Charles I ascended to the throne in 1625.[39] So, Sayeqa's evidence supports believing that Charles I

---

37 This case is first due to Radford (1966); what appears here is my own variant and development of the case. Radford uses the case to put pressure on the idea that knowledge entails belief, rather than to put pressure on KK. I am ambivalent as to whether the case fulfils this original ambition, though if it does, that may actually help my case (see fn. 40 below).

38 In this respect Sayeqa is not quite like the "clairvoyant" who is somehow reliable about some subject matter without possessing any identifiable evidence about the subject matter (other than her own direct insights of clairvoyants). What she lacks is decisive evidence *that* she has this evidence.

39 This assumes that one's evidence can support believing a falsehood. I take this to be the intuitive view—misleading evidence makes possible false but justified (in the sense of evidentially supported) beliefs. A few epistemologists (Sutton (2007); Littlejohn (2012)) challenge this orthodox view and contend that there are no justified false beliefs. Some also attribute this contention to Williamson, the leading opponent of the KK principle. But he himself explicitly assumes that justified false beliefs are possible (Williamson (2000: 9)). As he notes, this is consistent with his well-known claim that only knowledge can provide evidence.

ascended to the throne in 1625, but also supports believing that her evidence supports suspending judgment about (rather than believing) the proposition that Charles I ascended to the throne in 1625. So, Sayeqa is in a case of iterative failure.

Of course, someone might try to poke holes in these verdicts. For example, one might say that Sayeqa's feeling that her answer is a guess is enough to defeat the evidence she gained in the classroom at the first-order level, such that Sayeqa's evidence does not support believing that Charles I ascended to the throne in 1625. Or one might say (with considerably less plausibility, in my view) that Sayeqa's evidence gained in the classroom itself makes it the case that her evidence does not support believing that her evidence supports suspending judgment about whether Charles I ascended to the throne in 1625. It's hard to give knock-down responses to these claims, for the reasons given in the introduction to this section. In general, however, I do not see why these responses should be more convincing when it comes to *evidence* than when it comes to *knowledge*. And one can make the same protestations to deny that the case is an instance of KK failure: maintain that Sayeqa's sense that she is guessing defeats her first-order *knowledge*, or that her evidence from her classroom instruction secures higher-order knowledge that she knows that Charles I ascended to the throne in 1625. So if the case strikes us as a convincing instance of KK failure, I also think we should take it to be a convincing instance of iterative failure of evidence.

Someone might protest in reply that KK fails in the case for a reason that has nothing to do with the putative iterative failure of evidence. Specifically, the suggestion might be that KK fails in this case because Sayeqa does not even *believe* that she knows that Charles I ascended to the throne in 1625. It might then be claimed that Sayeqa is in a good enough evidential position to know this, if only she believed it. So the failure of KK is purely due to Sayeqa's not believing the higher-order proposition, and suggests no iterative failure of evidence.[40] Applied to all similar cases (as it would have to be to block (PIF)), the strategy would be to say that KK only *ever* fails for the (in one sense shallow) reason that belief sometimes fails to iterate.

I don't think this is right even in the case at hand. Given that Sayeqa can't *identify* that she is basing her belief on the evidence she got in the classroom (as opposed to random guessing), it does not seem to me that she is in a good evidential position to know that she knows that Charles I ascended to the throne. But even if one does find it plausible, we can complicate the case slightly so as to generate a plausible iterative failure of evidence even from

[40] Radford (1966), who first introduced the case, thought that Sayeqa does not even believe the *first-order* proposition in question, and that it was a case of knowledge without belief. If that's right, the KK failure cannot be due to a failure of belief to iterate, since belief is present at neither level.

an instance of KK failure that is only based on a failure of the iteration of belief. This exploits the role that knowledge plays in *providing* evidence.[41] Here the idea is that since one's knowledge plays a role in providing evidence for one's other beliefs, failures to know what one knows can lead to failures to know what one's evidence is, and so to iterative failure.

For this case, consider now the proposition that James I ceased to be King in 1625. Sayeqa knows, let's suppose, that James I immediately preceded Charles I, so that the year in which Charles I ascended to the throne would have to be the year in which James I ceased to be King. Given that Sayeqa knows that Charles I ascended to the throne in 1625, then, it seems right to say that Sayeqa's evidence supports believing that James I ceased to be King in 1625. However, the proposition that Charles I ascended to the throne in 1625 does not itself bear on the higher-order question of what doxastic attitude *Sayeqa's evidence* supports with respect to the proposition that James I ceased to be King in 1625. What would support believing the higher-order proposition that Sayeqa's evidence supports believing that James I ceased to be King in 1625 would be the corresponding higher-order proposition that Sayeqa's evidence includes the (first-order) proposition that Charles I ascended to the throne in 1625. But, since Sayeqa doesn't know that she knows that Charles I ascended to the throne in 1625, she is (plausibly) correspondingly unaware that this (first-order) proposition is part of her evidence. So the higher-order proposition isn't part of her evidence. So she lacks any significant evidence that her evidence supports believing that James I ceased to be King in 1625. So her evidence supports suspending judgment about whether her evidence supports believing that James I ceased to be King in 1625, even though her evidence supports believing that James I ceased to be King in 1625. So, again, we have a case of iterative failure.

The same method can be mimicked for generating cases of iterative failure out of failures of KK failure more generally. So if one accepts the possibility of KK failure, one should also accept the possibility of iterative failure due to being misled about what one's evidence is.

(b) Cases of being misled about the evidential support relations

Some (albeit a minority) of epistemologists deny the possibility of KK failure, and would presumably correspondingly deny the possibility of ignorance about what one's evidence is. But even if one is entirely correct about what one's evidence is, one can still be misled about which doxastic

[41] That knowledge plays such a role follows from Williamson's (2000: ch. 9) view that your evidence just is your knowledge, but also from various less bold claims. See Brown (2015) for a view of evidence that is critical of Williamson's whilst maintaining the core claim that your knowledge is generally part of your evidence.

attitudes this body of evidence supports, by being misled about the *support relations* that hold between items of evidence and doxastic attitudes. We turn now to iterative failures generated by this kind of misleading evidence.

I take the facts about the evidential support relations to be *normative* facts. As I said in section II, what the evidence supports can be understood in terms of what doxastic attitudes you have most reason to take given your evidence. As Thomas Kelly (2007) has convincingly argued, there is no plausible non-normative way of understanding the evidential support relation: it cannot be reduced to, say, statistical frequency. It must be understood as stating a normative claim about which propositions support which attitudes.

In general, it seems that one can often have misleading evidence about normative facts. A paradigm instance is when one receives misleading testimony. So in looking for a case of iterative failure, it is natural to look for a case where one receives misleading testimony about what one's evidence supports. Consider, then, the following case (loosely based on cases in Coates (2012) and Horowitz (2014)):

> **Miss Marple and Mabel.** Miss Marple is a detective who is famously good at assessing evidence. Miss Marple is investigating a murder that took place at the mansion on the hill, and she takes her great niece Mabel along with her. Miss Marple and Mabel set about the mansion collecting clues. Unfortunately, in their initial sweep of the house, nothing that they learn offers any kind of significant support to any particular hypothesis about who committed the crime. As part of her training of Mabel as her apprentice, after they have finished examining a crime scene, Miss Marple always tells Mabel what her own assessment of what the evidence supports. On this occasion, Miss Marple makes an uncharacteristic error, and declares to Mabel, "the clues lying around the house that you have seen up to this point support believing that the vicar did it".

What should we say about Mabel's position? By stipulation, the case is one where the clues lying around the house do not support believing that the vicar did it. However, Mabel has excellent reason to believe that what Miss Marple says about what the evidence supports is true. So it seems that this is a case whereby the clues lying around the house support suspending judgment about whether the vicar did it, but Mabel's evidence supports believing that the clues support believing that the vicar did it. The clues around the house are the crucial putative evidence that the vicar did it. And if the clues don't actually support believing that the vicar did it, Mabel's total evidence supports suspending judgment about whether the vicar did it. But Miss Marple's testimony is sufficient evidence to believe that the clues do support believing that the vicar did it, so Mabel's total evidence supports believing that her total evidence supports believing that the vicar it. That is a case of iterative failure.

Here, as with the examples in the last section, one can try to resist the case in a piecemeal way, exploiting the vagueness of how much evidential support is decisive. One can claim that Miss Marple's testimony is insufficient to support the higher-order belief that the evidence supports believing that the vicar did it. Or, going in the completely opposite direction, one can claim that Miss Marple's testimony is sufficient not just this higher-order belief, but also the *first-order* belief that the vicar did it. I'll begin, then, by offering some more theoretical reasons to think that this way of resisting the possibility of iterative failure won't work. This suggests that there will be a case structurally like that of Miss Marple and Mabel that amounts to iterative failure. The case of Miss Marple and Mabel is itself an illustration of the relevant structure. Unfortunately, it cannot be turned into a knockdown case simply by offering further details and precisification, due to the ineliminable vagueness of how much evidential support is decisive. But I hope to show that we have strong reasons to expect at least some cases of iterative failure that exploit the same structure. Then, I will consider some more general, principled arguments against such a possibility, and argue against them.

Let me start by giving what seems to me the right diagnosis of cases like that of Miss Marple and Mabel. It would be an overstatement, I agree, to say that Miss Marple's testimony about what the evidence supports *only* bears on the higher-order question of what the evidence supports. When Miss Marple testifies that the evidence supports believing that the vicar did it, she gives Mabel *some* pro tanto evidence in favor of the first-order claim that the vicar did it. In this way, higher-order evidence (at least usually) "trickles down" to support first-order attitudes: typically, "evidence of evidence is (some) evidence". However, I hold that the evidence provided by Miss Marple's testimony about what the evidence supports bears *less strongly* on attitudes towards the first-order question of whether the vicar did it than it does on attitudes towards the higher-order question of what the evidence supports.[42]

Some may find this immediately intuitive: after all, what Miss Marple testifies about directly is just the higher-order question, and the evidential impact of her testimony on the first-order question is seemingly more indirect. But here is a bit of theoretical machinery that will help to make good on this somewhat suggestive idea. All other things being equal, the strength of a piece of evidence in favor of believing some proposition will be mitigated by epistemic possibilities whereby that evidence obtains but the proposition in question is false. For example, suppose I tell you that the bus leaves at 3:30. Though this testimony provides good evidence in favor of believing that the bus leaves at 3:30, one thing that mitigates the strength of this evidence—that

[42] This moderate stance on higher-order evidence is also taken, as I read them, by Kelly (2010), Pryor (2013: 99–100), and Lasonen-Aarnio (ms.: 9).

prevents my testimony from supporting this belief to the maximum possible degree—is the (epistemic) possibility that I could be lying. By contrast, if you also know that I have taken a truth serum that prevents me from lying, the particular epistemic possibility that I am lying is eliminated. So all other things being equal, my testimony that the bus leaves at 3:30 is now (even) stronger evidence for your corresponding belief than it is in the more ordinary case. We can say that a particular *risk* (namely the epistemic possibility that I am lying) mitigates the strength of my evidence, so that when this risk is eliminated, my evidence is stronger.

Now, consider again the case of Miss Marple, who testifies that the evidence supports believing that the vicar did it. There are two beliefs that we want to consider here: the (higher-order) belief that the evidence supports believing that the vicar did it, and the (first-order) belief that the vicar did it.

When it comes to the strength of Miss Marple's testimony as it bears on the higher-order belief, things are relatively simple. The only epistemic possibilities that mitigate the strength of the testimony in favor of this proposition are those in which Miss Marple is speaking falsely (for whatever reason), since if what she says is true, then *ipso facto* the evidence supports believing the vicar did it, and so the higher-order proposition is true.

When it comes to the strength of Miss Marple's testimony as it bears on the first-order belief, however, we have more to contend with. As before, one thing that mitigates the strength of Miss Marple's testimony is the possibility that she is speaking falsely. In many of the possibilities in which Miss Marple speaks falsely, the vicar didn't do it. However, in contrast to the higher-order proposition, there are *also* epistemic possibilities in which Miss Marple is speaking truly, and yet the first-order proposition is still false. For it could be that the evidence *does* support believing that the vicar did it, but that the vicar nevertheless didn't do it; in such a case, the evidence on which she reports is misleading. So there are two sets of possibilities that mitigate the strength of Miss Marple's testimony in favor of the first-order proposition: the possibilities in which she speaks falsely, and the possibilities in which she speaks truly, but the evidence on which she reports is misleading and the vicar didn't in fact do it.

It is tempting at this point to argue as follows:

(1) Every possibility which mitigates the strength of Miss Marple's testimony as evidence for the higher-order belief also mitigates its strength as evidence for the first-order belief

(2) But, some possibilities which mitigate the strength of Miss Marple's testimony as evidence for the first-order belief do not mitigate its strength as evidence for the higher-order belief

(3) If (1) and (2), then all other things equal, Miss Marple's testimony is stronger evidence for the higher-order belief than it is for the first-order belief

(4) So (by (1), (2) and (3)), Miss Marple's testimony is stronger evidence for the higher-order belief than it is for the first-order belief

I think that this is on the right track, but as it stands it is too simple, for (1) is not quite right. There are some possibilities in which Miss Marple speaks falsely but the vicar *did* do it. In these possibilities, the higher-order belief is false but the first-order belief is true. Let us set out all the possibilities before us:

Set of possibilities	Fact of the matter as to what the evidence supports	Fact of the matter as to whether the vicar did it	Evidential support mitigated for which belief?
A	The evidence supports believing the vicar did it (Miss Marple speaks truly)	The vicar did it	Neither
B	The evidence supports believing the vicar did it (Miss Marple speaks truly)	The vicar didn't do it	First-order belief
C	The evidence does not support believing the vicar did it (Miss Marple speaks falsely)	The vicar did it	Higher-order belief
D	The evidence does not support believing the vicar did it (Miss Marple speaks falsely)	The vicar didn't do it	Both

Sets A, B, C and D jointly exhaust all the possibilities. We can ignore sets A and D, since these possibilities mitigate the support for neither and both of the beliefs, respectively. The real comparison we want is between sets B and C.

One might worry that sets B and C effectively "cancel each other out", so that the support provided by Miss Marple's testimony for the higher-order belief and the first-order belief ends up being equal (or, at least, that we have no general basis to assume that either belief will be more strongly supported). However, this would be too quick. For there are general grounds for holding that it is less likely that one of the C-possibilities obtains than that one of the B-possibilities obtains. In the C-possibilities, two unlikely things happen. First, Miss Marple speaks fal-

sely about a matter about which she is an expert. Second, and further, it turns out that, even though the evidence doesn't support believing that the vicar committed the crime (as she said it did), it also happens that he did commit it.[43] In the B-possibilities, by contrast, only one unlikely thing happens. Miss Marple speaks truly (as we would expect). But it turns out that even though the evidence supports believing that the vicar committed the crime, he actually didn't. So the B-possibilities involve fewer improbable events. This suggests that, absent any specific reasons to think otherwise, and certainly in at least *some* versions of the Miss Marple case, where her (known) reliability is fixed highly enough, the C-possibilities as a whole are antecedently less likely than the B-possibilities as a whole.[44] In other words, the possibilities that mitigate the strength of Miss Marple's testimony as evidence for the higher-order belief are more serious than those that mitigate the strength of Miss Marple's testimony as evidence for the first-order belief. This secures the result expressed in (4) above, even though the original simpler argument for it didn't quite succeed. Miss Marple's testimony is *better* evidence for the higher-order proposition about which she directly testifies than it is for the first-order proposition about who actually committed the crime.

By contrast, the clues themselves, if anything, bear more strongly on attitudes towards the first-order question of whether the vicar did it than they do on attitudes towards the higher-order question of what the evidence supports. Primarily, the clues speak to what attitude to take toward the first-order question of whether the vicar did it (in this case, they support suspending judgment). To the extent that the evidential import of the clues is (to some degree) self-evident, they also may be

---

[43] Admittedly, that is not *as* unlikely as the more specific scenario where the evidence *actively supports disbelieving* that the vicar committed the crime, and yet he did commit it. But it is still somewhat unlikely. Conditional only on the information that the evidence doesn't support believing that the vicar did it, it is still somewhat unlikely that the vicar did it. For this information rules out the possibilities in which the evidence supports believing that the vicar did it without ruling out the possibilities in which the evidence supports disbelieving that the vicar did it. Since, axiomatically, the majority of the possibilities in which the evidence supports believing that the vicar did it are ones in which the vicar did do it, ruling them out makes the proposition that the vicar did it less likely than it was antecedently.

[44] One might object that the first-order evidence (i.e. the clues observed in the hunt of the house) provide us with some reason to think that the C-possibilities rather than the B-possibilities obtain. But we are presently only talking about the evidential support afforded by Miss Marple's *testimony*, and its relative strength with respect to the first-order and higher-order beliefs. The clues may counterbalance against Miss Marple's testimony, but they do not affect the weight that it carries in and of itself.

able to support beliefs about their own nature, like the belief that they themselves support suspending judgment about whether the vicar did it. But this role seems secondary, and to require a greater strength of evidence to be decisive. In any event, the evidential weight of the clues at the higher-order level certainly isn't *stronger* than its weight at the first-order level.

To summarize, Miss Marple's testimony has more weight at the higher-order level than at the first-order level, while the clues around the house have more weight at the first-order level than at the higher-order level (or at least no more weight at the latter than the former). Putting these two points together, we get the result that, in some case at the margins, we should expect the clues themselves to outweigh Miss Marple's testimony at the level of her first-order attitude, but Miss Marple's testimony to outweigh the clues at the level of her higher-order attitude. That creates the possibility of iterative failure, where Mabel's (total) evidence supports the first-order attitude of suspension of judgment, but her (total) evidence also supports the second-higher attitude that her (total) evidence supports believing. Quibbling with individual cases by exploiting the vagueness of evidential support doesn't mitigate this possibility. However, there may be more systematic rationales for denying the possibility of iterative failure. I will now consider two such putative rationales.

The first strategy aims to establish that Miss Marple's testimony is not decisive even at the level of the higher-order attitude. The idea here is to claim that the justification of normative claims about evidential support relations is *a priori* and indefeasible by empirical testimony. Consequently, there are no false but evidentially supported beliefs about such claims. If that is right, no amount of misleading testimony from Miss Marple can outweigh Mabel's *a priori* justification for believing that the clues (given a fix on the empirical question of what the clues are) support suspending judgment.[45] This testimony is effectively inert, in the sense that it is said to be in principle incapable of tipping the balance in favor of believing that the clues support believing the vicar is guilty. What the evidence on balance supports with respect to what the clues support is fixed by one's *a priori*

[45] This traditionally unorthodox view that normative beliefs cannot be justified but false has recently received several defenses: from Titelbaum (2015), Whiting & Way (ms.) and Kiesewetter (ms.). Titelbaum restricts his claim of infallibility to beliefs about *rationality*, whereas the other two papers defend it with respect to all normative beliefs. If one distinguishes rationality from evidence-responsiveness, one can accept Titelbaum's claim (for which he has a powerful argument) without committing oneself to the impossibility of justified false beliefs about evidential support. Titelbaum himself seems to assume that his view does have this commitment, though he has elsewhere defended a very subjective notion of evidential support (Titelbaum (2010)), which may mitigate the radical-soundingness of his view.

justification only; Miss Marple's testimony is powerless to ever change what the evidence on balance supports with respect to this matter.

This, I think, is a very implausible verdict about the case. Mabel has excellent evidence that Miss Marple's ability to assess what the evidence supports is far superior to her own. Miss Marple is an expert—indeed, we can make her an arbitrarily reliable expert as long as we don't make her reliability *perfect*—about a particular subject matter: what the evidence supports. The idea that some person could be an expert about some subject matter about some matter to some arbitrarily reliable degree and yet their testimony about that subject matter be in principle incapable of affecting what the evidence on balance supports, is hard to stomach.

Note also that the principle is not just about misleading testimony, but rather blocks any other kind of empirical defeat of your *a priori* justification. So, for example,[46] suppose you discover that you are on a drug that makes people form inaccurate beliefs about what the evidence supports 99% of the time. As it happens you are in the lucky 1% who aren't affected. Your belief about what the evidence supports feels justified to you, but it would feel equally well-justified if you were on the drug and it were mistaken. In a case like this, the current strategy is committed to saying that your discovery that you are on the drug does nothing to undermine your belief about what the evidence supports.

More generally, consider the sorts of truths that the evidential support relations express. Remember that the position we are considering requires not just indefeasible *a priori* justification for *general* normative claims like 'believe what the evidence supports', but for the detailed normative facts of which pieces of evidence support which attitudes and to what degree. Such truths are, in my view, often radically unobvious to us. For example, what is the evidential support relation between the proposition that it seems to you that you have hands and the proposition that you have hands? What is the evidential support relation between the proposition that a theory is simple and the proposition that the theory is true? What is the evidential support relation between the proposition that past scientific theories failed and the proposition that our current theories are false? What is the evidential support relation between evidence that is very selectively filtered in biased ways and beliefs based on such evidence?

I think we are in pretty bad epistemic positions with respect to these questions—much worse epistemic positions than we face with respect to common empirical questions. Many people would have no idea how to answer them whatsoever. As such, the best evidence they have to go on about them will often be testimonial and by upbringing—by learning from others what to take as evidence for what. So I think it is highly implausible to say that there ques-

[46] This well-known sort of case originates from Christensen (2007: 10).

tions are ones about which everyone possesses indefeasible *a priori* evidence that renders (potentially misleading) empirical evidence irrelevant.

Moreover, even if one bites the bullet on these claims, this move will, I think, simply push the conflict between coherence and evidence-responsiveness elsewhere. Consider again the case of Mabel. Suppose that Mabel herself (correctly) judges Miss Marple to be much more reliable than she is in judging what the evidence supports. This itself carries certain commitments as to how to adjust her beliefs about what the evidence supports when she finds them to conflict with Miss Marple's. So, if she fails to treat Miss Marple's testimony as significant evidence with respect to the matter of what the evidence supports, without greatly downgrading Miss Marple's expert status (which is surely not warranted by one disagreement), she will be guilty of a different kind of incoherence. So, if one claims that the evidence supports Mabel maintaining her belief that the evidence supports suspending judgment in this case, even in the face of Miss Marple's testimony to the contrary, one purchases concord between evidence-responsiveness and (ILC) only at the price of conflict between evidence-responsiveness and other requirements. So the ultimate claim I am making that evidence-responsiveness and coherence requirements conflict remains plausible even in the face of this strategy.

So much for the first strategy. The second strategy, moving in the complete opposite direction, attempts to establish that if Miss Marple's testimony is decisive not only at the level of Mabel's higher-order attitude, but also at the level of her first-order attitude. Or, more cautiously, this strategy allows that Miss Marple's testimony *can* be decisive at the level of Mabel's higher-order attitude, and that *if* it is decisive at that level, it is also decisive at the level of Mabel's first-order attitude.

What kind of general principle could yield this result? One thought is to try to appeal to some general claim about (undercutting) defeat.[47] Undercutting defeaters, which are easiest to think of in the context of the attitude of belief, are supposed to be considerations that undermine the justification of a belief in a proposition p not necessarily by providing (sufficient) positive evidence to think that p is false, but rather merely by suggesting (perhaps misleadingly) that one's reasons for believing p are no good, in a way that neutralizes or mitigates their justificatory or evidential force.[48] The idea could be now that any justified belief to the effect that some doxastic attitude of one's own is not supported by the evidence is enough to serve as an undercutting defeater for that doxastic attitude, such that the attitude, ultimately, actually *isn't* supported by the evidence.

---

47 Cf., e.g., Bergmann (2005).

48 Cf., e.g., Pollock (1986: 38–39).

The problem with this principle is that it is far from clear how to extend it to the attitude of suspending judgment. The very idea of an undercutting defeater for suspending judgment is hard to understand. It would seem that the only way that suspension of judgment about p can be "defeated" is by one's gaining positive reasons to believe or disbelieve p. I deliberately designed the case of Miss Marple and Mabel as one in which it is suspending judgment that is the attitude that is being (putatively) defeated. Whereas the story about undercutting defeat might seem tempting for a case in which Miss Marple tells Mabel some *belief* of hers is not justified, it is does not extend to a case where Miss Marple tells Mabel that her *suspending judgment* is not justified.

So the opponent of iterative failure must hold that Mabel's attitude of suspending judgment is defeated not by being undercut but rather simply in virtue of Miss Marple's testimony giving her decisive, sufficient positive evidence for believing that the vicar did it. To achieve suitable generality in blocking iterative failure, the underlying claim here must be a strong principle about evidential strength: that if some piece of evidence is strong enough to be decisive at the level of the relevant higher-order attitude, then it is *always* also strong enough to be decisive at the level of the relevant first-order attitude. But, while I have conceded that Miss Marple's testimony has *some* evidential bearing on the first-order question of whether the vicar did it, I have already explained why we should reject the strong principle about the strength of this evidence. As I argued, Miss Marple's testimony bears *more strongly* on the second-order question of what the evidence supports than it does on the first-order question of who committed the crime. And vice-versa for the clues lying around the house. As a result, we should expect there to be a marginal case where Miss Marple's testimony outweighs the clues at the level of Mabel's second-order attitude, but the clues outweigh Miss Marple's testimony at the level of Mabel's first-order attitude. That will give us iterative failure, as well as falsifying the strong principle about evidential strength that was supposed to block it.

That brings me to the end of my defense of (PIF). I have argued that iterative failure can occur for two distinct reasons: either because one is misled about what one's evidence is (as in cases of KK failure such as that of Sayeqa, the unconfident examinee), or because one is misled about the evidential support relations (as in cases of misleading testimony such as that of Miss Marple and Mabel). If either contention is right, that is sufficient to establish (PIF).

V. Does (ILC) still have force in cases of iterative failure?

As I mentioned earlier, my argument would be in trouble if the rationale for (ILC) broke down whenever iterative failure obtained. So to complete my defense of (ILC), I want to briefly illustrate the way that it does not do so.

The crucial point here is that, at least in most cases, one cannot know that one is in a case of iterative failure. Let's illustrate this with respect to Sayeqa and Mabel. In Sayeqa's case, iterative failure is generated by a failure of KK. The proof that one cannot know that one is in a case of KK failure is simple.[49] Let p be the proposition that KK fails with respect to (in Sayeqa's case, that Charles I ascended to the throne in 1625), and let $K_S(p)$ be the claim that Sayeqa knows p. By the definition of KK failure:

(1) $K_S(p)$

(2) $\neg K_S(K_S(p))$

Suppose for reductio that Sayeqa knows that she is in a case of KK failure with respect to p, Sayeqa would have to know (1) and (2). In other words:

(3) $K_S(K_S(p))$

(4) $K_S(\neg K_S(K_S(p)))$

But (2) and (3) contradict each other. So, by reductio, Sayeqa cannot know that she is in a case of KK failure with respect to p.

Similar considerations apply to Mabel. For Mabel to know that she is in a case of iterative failure, she would have to know both that her evidence supports suspending judgment about whether the vicar did it and that her evidence supports believing that her evidence supports believing that the vicar did it. But if Mabel knew that her evidence actually supports suspending judgment about whether the vicar did it, she would be in a position to discount the misleading evidence that supports believing that her evidence supports believing that the vicar did it—in which case it wouldn't be misleading evidence at all, and she wouldn't be in a case of iterative failure.

The general lesson here is that iterative failure depends crucially on being misled in a way that precludes recognizing that you are in such a situation. This is significant. We might have a case for the failure of (ILC) if cases of iterative failure allowed one to provide a rationalization of why one is violating (ILC). We might imagine Mabel saying, "I see that normally I shouldn't violate (ILC), but this is a strange case where my evi-

[49] This point is also made by Williamson (2000: 119).

dence supports a particular doxastic attitude, but my evidence also supports believing that my evidence does not support this attitude". But actually, Mabel is not in a position to say this, for the reasons given above.

When one violates (ILC), one by definition *takes oneself* to be violating (ER). Imagine, then, what it would take to actually violate (ILC) (and respect (ER)) in a case of iterative failure. According to (ER), Mabel is rationally required (i) to suspend judgment about whether the vicar did it and (ii) to believe that her evidence supports believing that the vicar did it. Suppose that Mabel actually does this, as (ER) demands of her, and that she is aware of her own doxastic states. In that case, Mabel *takes* herself to be violating (ER): as far as she is concerned, this is just any old case where she's failed to believe what her evidence supports. But since she is in a case of iterative failure, it turns out that by luck she does satisfy it, in spite of herself.

It's an interesting question whether one can ever be in the position that (ER) diagnoses this case as putting Mabel in—namely, where the only way to be rational is by luck, in spite of believing oneself to be failing.[50] If one cannot be in such a position, this presents a further reason to suspect that (ER) is not best understood as a requirement of *rationality*. Regardless of how we answer this question, however, the point I want to make is that this kind of lucky satisfaction of (ER) certainly does not seem to diminish the distinctive incoherence of her violating (ILC). Mabel can give no rationalization for her violation of (ILC) in terms of her need to satisfy (ER), since by her lights she thinks she is violating (ER). So the incoherence of her (ILC)-violating mental states is just as clear in this case as it is with any other violation of (ILC). And the argument for (ILC) given earlier in terms of the relationship of a judgment of evidence to a judgment of likelihood is unaffected.

VI. What kind of conflict?

Let's take stock. I've argued that cases of iterative failure are possible, but that (ILC) is nevertheless a genuine coherence requirement. This shows that evidence-responsiveness (responsiveness to epistemic or evidential reasons) and conformity to coherence requirements can come into conflict with one another: there are cases where one cannot have both. So Superfluousness is false. We cannot say that rationality consists solely in evidence-responsiveness, and expect that to capture our intuitions about the irrationality of certain forms of incoherence derivatively. That provides us with reason to reject Evidentialism, the view that one is

[50] I do genuinely mean that it is interesting. I am not coyly expressing contempt for the view that such a case is possible. For some considerations that may support believing that it is, see Hawthorne & Srinivasan (2013); Srinivasan (forthcoming).

rational iff one satisfies (ER). It also shows that coherence requirements will have to be theorized, at least to some degree, in their own right, in the sense that they will not just fall out of our best account of evidential reasons.

As I mentioned in section I, there are two ways of understanding this conflict between coherence requirements and evidence-responsiveness. One is to say that both coherence requirements like (ILC) on one hand, and (ER) on the other, are genuine rational requirements, in some recognizable and unified sense of the term 'rational'. Then, cases of iterative failure will count as rational dilemmas—cases where whatever one does, one is not rational. This first option is, on one reading, endorsed by David Christensen (2007, 2010, and esp. 2013: 92–96). Christensen's view is that both (ER) and (ILC) are "rational ideals", and that in cases of iterative failure, one "will end up violating some ideal or other". The second possibility is to say that really there are two fundamentally different kinds of normative demand here, that need to be pulled apart. In that case, we have conflicts *between* two different normative domains, but not *within* one single normative domain. This is the view I will be defending.

Lest the distinction between these views seem somewhat murky, let me say a little more about it by way of an analogy. Sometimes, one's moral reasons point one way while one's prudential reasons point another. Now, in a weak sense, this is just obviously a conflict between normative demands of two different kinds—one set of reasons is *moral*, while another set is *prudential*. However, on what I take to the majority (but not hegemonic) view in ethics, while moral and prudential reasons are trivially different in the sense that they are reasons issuing from different kinds of consideration, they are still ultimately commensurable. Here are some (putative) markers of this commensurability. (1) They *weigh against* one another, and jointly contribute to determining the answer to a broader question, that of what one has most reason, all-things-considered, to do. So, when the two come into conflict, they *compete*, and one "wins". (2) They enter into deliberation in the same sorts of ways, and it makes sense in deliberation to think "one the one hand, there are these moral considerations, and on the other hand, there are these prudential considerations. . ." (3) They both can be cashed out in terms of a unified notion of a reason, understood in terms of a fact counting in favor of something. (4) They both count in favor of the same *sorts* of things: individual actions. In this way, although moral and prudential reasons differ trivially, they are still demands of a same, fundamental, broader normative kind.

However, in various ways and to varying degrees, one could imagine someone questioning these assumptions, and holding that moral reasons

and prudential reasons are in fact *not* commensurable, in one or more of the above senses. It might be held that there is no such thing as what one "all-things-considered" has most reason to do; only what one has most moral reason to do and what one has most prudential reason to do. (This may have been Sidgwick's view.) Similarly, one might hold that they somehow enter deliberation in different ways, or that they are not both reasons in a single, unified sense of a 'reason'. Perhaps they even don't count in favor of the same sorts of things (for example, on some views, prudential reasons might be reasons for individual actions, whereas moral reasons might primarily apply to intentions, habits, rules, or maxims). To the extent that one made these claims, one would be holding that moral and prudential reasons were normative demands of fundamentally different kinds, and not just in the trivial sense that the former are moral and the latter prudential.

I take no stand here on the right account of the depth of difference between moral and prudential reasons. I introduce the analogy only to illustrate the difference between holding that evidential reasons and coherence requirements are demands of the same fundamental kind, and that of holding that they are demands of fundamentally different kinds. The most natural reading of Christensen, I think, suggests a view broadly of the first sort. He does not distinguish coherence requirements and evidential reasons explicitly as different kinds of normative requirement. While he could probably recognize some distinction and allow that they differ in a trivial sense, he seems to think of them both as demands of some broader normative kind, and at times even writes as if it is a *further* question, given a "conflict of ideals", which ideal it would be rational to satisfy; as if the two are merely competing *pro tanto* forces.[51]

By contrast, I prefer the second view. I will give three broad sets of considerations in favor of preferring the second view to the first. The first points to a number of interrelated metaphysical differences between coherence requirements and evidential reasons. To criticize someone as incoherent is to accuse them in some way of failing by their own standards. In such a case, their attitudes fail to fit together in the right way: they inhabit a perspective that does not make sense from a first-person perspective. One can make such a criticism without taking any stand on what such a subject's evidence actually supports. Such a criticism is in one important way less substantive—yet often more demonstrable—than the charge that the subject has failed to respond to their evidence correctly. This difference is reflected in the difference in scope that I drew attention to at the start of section III. Coherence requirements are wide-scope, and do not speak in favour in individual attitudes *simpliciter* but

[51] See, e.g., Christensen 2013: 96.

rather against particular combinations of attitudes. Evidential reasons, by contrast, are narrow-scope: when you have a particular body of evidence, it speaks in favor of some determinate attitude or attitudes. From your having the evidence, we can derive the conclusion that you have reason to have those attitudes, simpliciter.

Relatedly, coherence requirements and evidential reasons play very different sorts of roles in how they guide reasoning. As Niko Kolodny (2005: see esp. 547) has argued, one is not usually guided in deliberation about what to believe by the thought that by believing something, one can avoid incoherence of the sort banned by (ILC). Rather, one is guided by the thought that there is weighty evidence that the proposition to be believed is true. Nevertheless, this process is in one sense *regulated* by one's background disposition to satisfy (ILC)—that is, to take the doxastic attitudes that one *takes oneself* to have conclusive evidence for, and to refrain from taking those attitudes that one takes oneself not to have conclusive evidence for. This difference in the role that coherence requirements and evidential reasons play in reasoning reinforces their metaphysical distinctness. Though I have argued that they can *conflict* in the sense that sometimes one cannot satisfy both, they do not *compete* as *pro tanto* forces at the same level of one's deliberations, in answering some broader question of what one ought to believe *simpliciter*.

This point is reinforced further by the related fact that when one satisfies (ILC), it will *seem* to one that one is satisfying (ER), and when one violates (ILC), it will seem to one that one is violating (ER). Consequently, as I already argued in section V, it is normally not possible for one to know that one is in a situation of iterative failure, such that the two conflict. These considerations all taken together make it, in my view, implausible to think of (ER) and (ILC) as, in Christensen's language, "competing ideals" of the same kind which one has to somehow weigh against each other in deliberation.

The second argument in favor of the second approach builds on the first. Suppose I am right that coherence requirements an evidential reasons are not competing *pro tanto* forces, such that they both contribute to some broader determination of what one ought to believe *simpliciter*. Then, an approach which says that both normative demands are to be understood as requirements of rationality will lead to the positing of rational dilemmas, whereby an agent is irrational *whatever* she does. This might seem like hair-splitting, since on the second approach we still have to acknowledge conflicts *between* different normative domains, such that agents must fall short of normative ideality in *some* respect. I'm not sure that this is as bad a thing to have to say: there does seem to be something about *irrationality* in particular that should be in some way due to the agent herself, and not simply due to a situation outside her control.

In any case, though, there are formal advantages to representing normative conflicts as occurring between different normative domains, rather than within a single domain. For example, it allows us to preserve the axioms of standard deontic logic. For example, if you are rationally required to Φ and rationally required not to Φ, then by standard deontic logic you are rationally required (to Φ and not to Φ). In other words, you are rationally required to do the impossible. If, conversely, the normative requirement to Φ is of a fundamentally different kind to the normative requirement not to Φ, then there is no *single* sense in which you are required, even by the lights of standard deontic logic, (to Φ and not to Φ).

Furthermore, allowing *just any* kind of conflict between requirements is permissive in a way that makes the methodology of arguing for particular requirements considerably more difficult. One check on our ability to posit rational requirements costlessly is the possibility that such rational requirements might conflict with other, more plausible, requirements. If we allow for rational dilemmas, then we can never show a putative requirement of rationality to be false by showing that it conflicts with some other important requirement. By contrast, if we allow conflicts across normative domains but not within them, the ban on intra-domain conflicts will still give us some way to rule at least some putative requirements out.

Last, as the puzzle that we have been considering throughout illustrates how confusing coherence requirements with claims about reasons is liable to lead to substantive philosophical mistakes. Some have thought that (ILC) can be simply be restated as a claim about reasons:

> **Inter-level reasons (ILR)**
> (i) If S has most epistemic reason to believe that her evidence supports D (p), then S has most epistemic reason to take D(p).
> (ii) If S has most epistemic reason to believe that her evidence does not support D(p), then S has most epistemic reason not to take D(p).[52]

Given the equivalence of what one's evidence (*pro toto*) supports believing to what one has most epistemic reason to believe, the second part of (ILR) just is the denial of (PIF). So if one takes the inter-level principle to statable as a claim about reasons, it is clear why one would think that the inter-level principle is incompatible with (PIF), such that one has to choose between them, and such that the acceptance of one provides conclusive grounds for the rejection of the other. But although either (ILC) or (ILR) might be picked out by the vague term "level-bridging principles",

[52] See, e.g., White (2007: 120). Dan Greco also suggested this to me in correspondence.

the two claims are really of fundamentally different kinds. Whereas (ILC) is a wide-scope coherence requirement, (ILR) is not, strictly speaking, a requirement at all, but rather a strong metaphysical constraint on the reasons that an agent it is possible for an agent to have at any given point in time. Specifically, (ILR) posits the inability to be evidentially misled about your own epistemic reasons in a systematic way. This is simply not a commitment of (ILC).

If a philosopher confuses (ILC) with (ILR), and moves between them uncarefully, she will be lead to the substantively mistaken philosophical view that (ILC) and (PIF) are incompatible claims. This is what has, I think, lead so many philosophers to argue from one to the rejection of the other. The importance of avoiding this mistake and others like it provides us with another good reason to sharply distinguish conformity to coherence requirements from evidence-responsiveness, and not to place them under a single undifferentiated heading.

VII. Conclusion

As I said in the Introduction, my own view is that the clearest way to mark this important difference between coherence requirements and claims about evidence-responsiveness is to follow many theorists of practical rationality in restricting the use of the term 'rationality' to the satisfaction of coherence requirements. We would then replace (ER), a claim about rationality, with (ER*), a claim about reasons. I do not insist on this way of speaking. Suppose that we do follow it, however. Once this stipulation about 'rationality' is in place, no doubt many epistemologists would recognize that (ER) is to be rejected. Of course, they will say, in the sense of 'rationality' that concerns coherence alone, rationality does not require responding to one's evidence. So is my dispute with epistemological orthodoxy merely terminological?

No. In assuming that one of (ILC) and (PIF) has to go in response to our puzzle, epistemologists have been tacitly taking it for granted that the correct formulations of (ER) and (ILC) make use of the same normative concept. Otherwise there is simply no incompatibility between (ILC) and (PIF). So it will not do for epistemologists to simply insist that all along, they only intended (ER) to be a claim about reasons, and not about rationality *in my sense*. Drawing this distinction does not leave everything as it was.

The view we have arrived at, by contrast, is that the puzzle introduced in section I is ultimately to be (dis)solved by distinguishing evidence-responsiveness on one hand, and coherence on the other. However we use terms, (ER) and (ILC) are not to be stated using the same normative concept. In cases of iterative failure, we will say, there is simply a conflict between responding to one's reasons, and being coherent ("rational", if we are feeling brave). Some may find that unsatisfying: "but what *ought* I to do in

such a case?", they will ask. If this 'ought' is the 'ought' of 'having most reason', the answer is clear: you ought to believe what your evidence supports. Unfortunately, the tragedy of a situation where you have all-things-misleading higher-order evidence is that you cannot do this (at every level of your beliefs) in a way that maintains coherence. That is what this kind of epistemic ignorance about what your own evidence supports consigns you to. That's the bad news.

The good news, though, is that this kind of agonizing conflict won't be tearing you apart any time soon. For, as I noted in section VI above, it is also part of the misleadingness of misleading higher-order evidence that you cannot know it is misleading. As such, when conflicts between satisfying (ILC) and responding to your evidence arise, you will not know that they have arisen. If you are rational, you will go on satisfying (ILC), and in virtue of that, thinking that you have responded to your evidence correctly. From your perspective, there will be no deliberative dilemma about whether to satisfy (ILC) or whether to take the attitudes that your evidence supports. That's the good news. But you will be wrong: unbeknownst to you, you will not have taken (all) of the attitudes that your evidence supports. That's the bad news. . .you get the idea.

I hope that this paper has, among other things, served as a case study in how coherence requirements, as contrasted with claims about evidence-responsiveness, are of importance to epistemology. My own view is that several other epistemological debates—about the relationship between belief and credence, about putative requirements of deductive consistency, about correct inference, about belief updating, and about peer disagreement—have been distorted by attempts to formulate central claims and principles in terms of evidential reasons rather than coherence requirements. Showing this is a task for (a number of) other days.[53]

[53] For helpful discussions related to this paper I'm grateful to Facundo Alonso, Caroline Arruda, Samuel Asarnow, George Bealer, Julia Borcherding, John Broome, Jennifer Carr, Julianne Chung, Stephen Darwall, Keith DeRose, Daniel Fogal, Tamar Gendler, Matt Kotzen, Josh Knobe, Clayton Littlejohn, Jennifer Morton, Jessie Munton, Ram Neta, Kate Nolfi, John Pittard, David Plunkett, Jim Pryor, Andrew Reisner, Geoff Sayre-McCord, Sun-Joo Shin, Sam Shpall, Zoltán Szabó, Bruno Whittle, Daniel Wodak, and audiences at LSE, Yale, and the St Louis Annual Conference on Reasons and Rationality. I'm especially indebted to Maria Lasonen-Aarnio, Ralph Wedgwood, Daniel Whiting, several anonymous referees (including an especially helpful referee for *PPR*), and above all Dan Greco for detailed comments on earlier written drafts. Since I wrote the original draft of this paper, a number of others have written manuscripts on similar puzzles, including Littlejohn (forthcoming), Pryor (ms.), and especially Lasonen-Aarnio (ms.). Although we all initially developed our papers independently, and each of us resolves the relevant puzzle in different ways, I have very much benefitted from reading their work in revising the present paper.

References

Adler, J.E. (2002). *Belief's Own Ethics*. Cambridge, MA: MIT Press.

Ballantyne, N. & Coffman, E.J. (2011). 'Uniqueness, Evidence, and Rationality,' *Philosophers' Imprint*, 11/18.

Bergmann, M. 2005. 'Defeaters and Higher-Level Requirements,' *Philosophical Quarterly*, 55/220: 419–436.

BonJour, L. (1985). *The Structure of Empirical Knowledge*. Cambridge, MA: Harvard University Press.

Broome, J. (1999). 'Normative Requirements,' *Ratio* 12: 398–419.

—— (2013). *Rationality Through Reasoning*. Chichester: Wiley-Blackwell.

Brown, J. (2015). 'Evidence and Epistemic Evaluation,' *Oxford Studies in Epistemology*, 5.

Brunero, J. (2010). 'The Scope of Rational Requirements,' *Philosophical Quarterly*, 60/238: 28–49.

Christensen, D. (2007). 'Does Murphy's Law Apply in Epistemology? Self-Doubt and Rational Ideals,'*Oxford Studies in Epistemology*, 2: 3–31.

—— (2010). 'Higher-Order Evidence,' *Philosophy and Phenomenological Research*, 81/1: 185–215.

—— (2013). 'Epistemic Modesty Defended,' in Christensen & Lackey (2013).

—— & Lackey, J. (eds.) (2013). *The Epistemology of Disagreement: New Essays*. Oxford: Oxford University Press.

Coates, A. (2012). 'Rational Epistemic Akrasia,' *American Philosophical Quarterly*, 49/2: 113–124.

Conee, E. & Feldman, R. (1985). 'Evidentialism,' *Philosophical Studies*, 48/1: 15–34.

Dancy, J. (2000). *Practical Reality*. Oxford: Oxford University Press.

Easwaran, K. & Fitelson, B. (2015). 'Accuracy, Coherence, and Evidence,' *Oxford Studies in Epistemology*, 5.

Elga, A. (2005). 'On Overrating Oneself. . .and Knowing It,' *Philosophical Studies*, 133/1-2: 115–124.

Fantl, J. & McGrath, M. (2009). *Knowledge in an Uncertain World*. Oxford: Oxford University Press.

Feldman, R. (1988). 'Subjective and Objective Justification in Ethics and Epistemology,' *The Monist*, 71/3: 405–419.

—— (2005). 'Respecting the Evidence,' *Philosophical Perspectives* 19: 95–119.

Fogal, D. (ms.). 'Rationality and Requirements,' draft manuscript, New York University.

Foley, R. (1993). *Working without a Net*. Oxford: Oxford University Press.

Goldman, A. (1986). *Epistemology and Cognition*. Cambridge, MA: Harvard University Press.

—— (2009). 'Internalism, Externalism, and the Architecture of Justification,' *Journal of Philosophy*, 106/6: 309–338.
Greco, D. (2014). 'A Puzzle about Epistemic Akrasia,' *Philosophical Studies*, 167/2: 201–219.
Harman, G. (1986). *Change in View*. Cambridge, MA: MIT Press.
Hawthorne, J. & Srinivasan, A. (2013). 'Disagreement Without Transparency: Some Bleak Thoughts,' in Christensen & Lackey (2013).
Hazlett, A. (2012). 'Higher-Order Epistemic Attitudes and Intellectual Humility,' *Episteme*, 9/3: 205–223.
Horowitz, S. (2014). 'Epistemic Akrasia,' *Noûs*, 48/4: 718–744.
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: 163–96.
—— (2006). 'Evidence,' *Stanford Encyclopedia of Philosophy*.
—— (2007). 'Evidence and Normativity: Reply to Leite,' *Philosophy and Phenomenological Research*, 75/2: 465–474.
—— (2010). 'Peer Disagreement and Higher Order Evidence,' in Feldman & Warfield (eds.), *Disagreement*. Oxford: Oxford University Press.
Kiesewetter, B. (ms.). 'You Ought to Φ Only If You May Believe That You Ought to Φ,' draft manuscript, Humboldt University Berlin.
Kolodny, N. (2005). 'Why be Rational?,' *Mind*, 114/455: 509–563.
—— (2007). 'How does Coherence Matter?,' *Proceedings of the Aristotelian Society*, 107/3: 229–263.
Lasonen-Aarnio, M. (ms.). 'Enkrasia or Evidentialism? Learning to Love Mismatch,' draft manuscript, University of Michigan.
Littlejohn, C. (2012). *Justification and the Truth-Connection*. Cambridge, UK: Cambridge University Press.
—— (forthcoming). 'Stop Making Sense? On a Puzzle about Rationality,' *Philosophy and Phenomenological Research*.
Lord, E. (2014). 'The Coherent and the Rational,' *Analytic Philosophy*, Online Early View.
McDowell, J. (1998). *Mind, Value & Reality*. Cambridge, MA: Harvard University Press.
Owens, D. (2000). *Reason Without Freedom*. London: Routledge.
Parfit, D. (2011). *On What Matters: Vol. 1*. Oxford: Oxford University Press.
Pollock, J. (1986). *Contemporary Theories of Knowledge*. Savage: Rowman & Littlefield.
Pritchard, D. (2002). 'Radical Scepticism, Epistemological Externalism, and Closure,' *Theoria*, 68/2: 129–161.

Pryor, J. (2004). 'What's Wrong with Moore's Argument?,' *Philosophical Issues*, 14: 349–378.
—— (2013). 'Problems for Credulism,' in Tucker (ed.), *Seemings and Justification*. Oxford: Oxford University Press.
—— (ms.). 'The Merits of Incoherence,' draft manuscript, New York University.
Radford, C. (1966). 'Knowledge—By Examples,' *Analysis*, 27: 1–11.
Rawls, J. (1996). *Political Liberalism*. New York: Columbia University Press.
Raz, J. (2011). *From Normativity to Responsibility*. Oxford: Oxford University Press.
Ridge, M. (2014). *Impassioned Belief*. Oxford: Oxford University Press.
Scanlon, T.M. (1998). *What We Owe to Each Other*. Cambridge, MA: Harvard University Press.
—— (2007). 'Structural Irrationality,' in Brennan, Goodin, Jackson & Smith (eds.), *Common Minds: Themes from the Philosophy of Philip Pettit*. Oxford: Oxford University Press.
Schroeder, M. (2009). 'Means-end Coherence, Stringency, and Subjective Reasons,' *Philosophical Studies*, 143: 223–248.
—— (2012). 'The Ubiquity of State-Given Reasons,' *Ethics*, 122/3: 457–88.
Schwitzgebel, E. (2008). 'The Unreliability of Naïve Introspection,' *Philosophical Review*, 117/2: 245–73.
Setiya, K. (2004). 'Against Internalism,' *Noûs*, 38/2: 266–298.
Smith, M. (2007). 'Is There A Nexus Between Reasons and Rationality?,' in Tenenbaum (ed.), *Moral Psychology*. Amsterdam: Rodopi.
Smithies, D. (2012). 'Moore's Paradox and the Accessibility of Justification,' *Philosophy and Phenomenological Research*, 85/2: 273–300.
Srinivasan, A. (2013). 'Are We Luminous?,' *Philosophy and Phenomenological Research*, Online Early View.
—— (forthcoming). 'Normativity without Cartesian Privilege,' *Philosophical Issues*.
Sutton, J. (2007). *Without Justification*. Cambridge, MA: MIT Press.
Sylvan, K. (ms.). 'On Divorcing the Rational and the Justified in Epistemology,' draft manuscript, University of Southampton.
Titelbaum, M.G. (2010). 'Not Enough There There: Evidence, Reasons, and Language Independence,' *Philosophical Perspectives*, 24: 477–528.
—— (2015). 'Rationality's Fixed Point (or: In Defense of Right Reason),' *Oxford Studies in Epistemology*, 5.
Wallace, R.J. (2001). 'Normativity, Commitment, and Instrumental Reason,' *Philosophers' Imprint*, 1/4.
Way, J. (2010). 'Defending the Wide-Scope Approach to Instrumental Reason,' *Philosophical Studies*, 147/2: 213–233.

Weatherson, B. (ms.) 'Do Judgments Screen Evidence?,' draft manuscript, University of Michigan.
Wedgwood, R. (2012). 'Justified Inference,' *Synthese*, 189/2: 273–295.
White, R. (2005). 'Epistemic Permissiveness,' *Philosophical Perspectives*, 19: 445–459.
—— (2007). 'Epistemic Subjectivism,' *Episteme*, 4/1: 115–129.
Whiting, D. and Way, J (ms.). 'If You Justifiably Believe That You Ought to F, You Ought to F,' draft manuscript, University of Southampton.
Williamson, T. (2000). *Knowledge and its Limits*. Oxford: Oxford University Press.
Worsnip, A. (forthcoming). 'Moral Reasons, Epistemic Reasons, and Rationality,' *Philosophical Quarterly*. doi: 10.1093/pq/pqv084