

Evidence-Coherence Conflicts Revisited

Alex Worsnip

Forthcoming in Nick Hughes (ed.), *Epistemic Dilemmas*, Oxford University Press
Penultimate draft; please cite published version where possible

There are at least two different aspects of our rational evaluation of agents' doxastic attitudes (and absences thereof). First, we evaluate these attitudes according to whether they are supported by one's evidence. It's this kind of rational evaluation that's (usually) at issue when we say that it is irrational to deny anthropogenic climate change, or to believe in a flat earth, or to believe that there are fairies living at the bottom of your garden. Second, we evaluate these attitudes according to how well they cohere with one another. It's this kind of rational evaluation that's (usually) at issue when we say that it's irrational to have contradictory or otherwise inconsistent beliefs, or that it's irrational to have credences that violate the probability axioms, or that it's irrational to believe something that *you yourself* judge yourself not to have good evidence for.

The distinction between these two kinds of rational evaluation generalizes beyond doxastic attitudes. Assuming that evidential considerations constitute at least one very prominent kind of *reason* for and against doxastic attitudes, failures to correctly respond to one's evidence are ipso facto failures to respond correctly to one's reasons - just as one can fail to respond correctly to one's reasons in one's intentions, preferences, hopes, fears, and so on. Call rational evaluations that concern the responsiveness of one's attitudes to one's reasons evaluations of *substantive* rationality. Likewise, there are many forms of incoherence that involve non-doxastic attitudes: cyclical preferences, inconsistent intentions, failures to intend to believed means to one's ends, and the like. Call rational evaluations that concern the coherence of one's attitudes evaluations of *structural* rationality.

A central question in the theory of rationality is this: are substantive and structural rationality really two distinct, *sui generis*, equally genuine kinds of rationality? Or can one of them be eliminated or reduced to (or even shown to be identical to) the other? Call those who take the former view *dualists* about rationality, and those who take the latter view *monists* about rationality.

There are many different fronts on which the battle between dualism and monism can be waged. But one way to argue for dualism – suggested by my (2018) – is to argue that the requirements of substantive and structural rationality can in fact come into conflict. As I'll explain shortly, if this is true, a prominent monist strategy for eliminating *sui generis* requirements of structural rationality immediately fails.

In this paper, I want to push the dialectic on this issue further forward in two main ways. In the first main part of the paper (§§4-5), I'll revisit conflicts between substantive rationality (reasons-responsiveness; or, in the doxastic case, evidence-responsiveness) and structural rationality (coherence) in more depth. I will discuss two kinds of cases – one that my (2018) discussed (cases of misleading higher-order evidence), and one that it didn't (preface cases). I will argue that both cases can plausibly, but not utterly conclusively, be diagnosed as evidence-coherence (i.e. substantive-structural) conflicts. But I'll also argue – and this goes beyond what I contended even about misleading

higher-order evidence cases in my (2018) – that the most plausible ways of *resisting* the diagnosis of these cases as conflict cases still end up diagnosing them in ways that undermines monism.

In the second main part of the paper (§§6-7), I'll suppose for the sake of argument that we *should* understand the relevant cases as conflict cases, and try to say something about what we should do if and when such conflicts arise. I demurred from doing this in my (2018), thinking that there is nothing to be said about how such conflicts should be adjudicated. They are, I thought, simply dilemmatic, in the strong sense that there is no all-things-considered answer as to what one should do in such cases. However, I've changed my mind about this, at least in the case of conflicts generated by misleading higher-order evidence. I will argue that in those cases, the coherence (structural) requirements take precedence over the evidential (substantive) requirements. Unfortunately, however (and this is a matter of genuine disappointment to me), the arguments I give for this don't seem to generalize to preface cases. So *if* preface cases are conflicts of evidence and coherence, I don't (currently) see a way to resist understanding them as dilemmatic in a strong sense.

First, though, a bit more setup.

1. A slightly more precise fix on substantive rationality

In introducing the distinction between substantive and structural rationality, I said that substantive rationality concerns responding correctly to our reasons. But this is somewhat imprecise as it stands. It needs two qualifiers.

First, substantive rationality does not require responsiveness to reasons that are completely outside of one's epistemic ken. Suppose that \$1,000,000 is buried in John's back yard, but you don't know this and you have no evidence whatsoever that indicates it. There's perhaps a sense in which the fact that \$1,000,000 is buried in your back yard is nevertheless a strong reason to (intend to) dig up your back yard. But you would, I take it, in no ordinary sense be *irrational* for failing to do so, given that you're in no epistemic position to appreciate the relevant reason. Similarly, it may make sense to talk of reasons for belief that are completely outside of one's epistemic ken. Suppose that you're investigating Paul's murder. Suppose further that Ringo stands to inherit all of Paul's money if Paul dies, but you don't know this and you have no evidence whatsoever that indicates it. There's perhaps a sense in which the fact that Ringo stands to inherit all of Paul's money if Paul dies is a strong reason to believe that Ringo killed Paul. But you would, I take it, in no ordinary sense be irrational to failing to believe this, given that you're in no epistemic position to appreciate the relevant reason.

So substantive rationality requires you only to respond to the reasons that are, loosely speaking, within your epistemic ken. We might express this by saying that substantive rationality requires you only to respond to reasons that you "possess",¹ or that it requires you only to respond to your "evidence-relative" reasons.² In the doxastic case, the slogan that (substantive) rationality requires you to respond to your evidence already reflects this: it restricts itself to the reasons for belief provided by the evidence that you have.

¹ Lord 2018.

² Kiesewetter 2017.

Some has thought that once we see that rationality requires you only to respond to the reasons within your epistemic ken, the distinction between substantive and structural rationality melts away.³ But this is too hasty. Consider the denier of anthropogenic climate change. Let's stipulate that this person is aware of what is, in fact, decisive evidence against her belief that anthropogenic climate change is not real (say, the testimony of experts to the effect that anthropogenic climate change is real). In that case, she believes something that her evidence decisively tells against – she is substantively irrational. Nevertheless, she might be perfectly structurally rational – her beliefs might be entirely consistent, her credences might obey the probability axioms, she might not be defying her own *judgments* about what her evidence supports, and so on.⁴

Second, substantive rationality does not require responsiveness to so-called “wrong-kind” reasons. The paradigm of a wrong-kind reason for believing p is a financial offer of money to believe p . Again, perhaps wrong-kind reasons are genuine normative reasons; I need not take a stand on this. But if they are, they don't bear on the *rationality* (substantive or otherwise) of believing p .⁵ Intuitively, the fact that you've bribed me \$1,000,000 to believe that $2+2=5$ doesn't make the belief any less irrational; at most it gives me a reason to be irrational. Moreover, wrong-kind reasons are notoriously difficult to respond to: it is very hard for me to believe that $2+2=5$ on the basis of your bribing me \$1,000,000 to do so. In light of this, it seems very harsh to say that I'm *irrational* if I don't believe that $2+2=5$. The point generalizes the wrong-kind reasons for other attitudes such as intention and preference. Again, the slogan that (substantive) rationality requires you to respond to your evidence already reflects this, since evidence is the paradigmatic *right-kind* reason for belief.

Note that if substantive rationality *did* require responsiveness to wrong-kind reasons, then it would be extremely easy to generate conflicts between substantive and structural rationality. We could just dream up a case where you're offered tons of money to violate a structural requirement. But I assume that conflicts between substantive and structural rationality doesn't come this cheaply.

Now we can state the claim that substantive rationality requires you to respond correctly to your reasons a bit more precisely. Substantive rationality requires you to respond correctly to your *right-kind, possessed* (or evidence-relative) reasons.

2. How substantive-structural conflicts strengthen the case for dualism

Prima facie, substantive and structural rationality seem distinct, and they both seem to pick up on genuine features of our talk about rationality. How, then, might one try to be a monist about rationality? A recently prominent monist proposal builds on the following hypothesis:

³ Wedgwood 2017.

⁴ It might now be claimed that since failures to respond correctly to one's evidence themselves involve a bad of “bad relation” between one's mental states, they are themselves a kind of incoherence. Perhaps there's a sense of ‘incoherence’ in which this is true, but I think it's not the (narrower) sense of ‘incoherence’ that grounds structural irrationality. See Worsnip (2021: chapter 1) for more.

⁵ Cf. Kelly 2002; Schroeder 2012.

The Guarantee Hypothesis. For any set of attitudinal mental states⁶ $\{A_1 \dots A_n\}$ of the kind associated with structural irrationality, it is guaranteed that at least one of $\{A_1 \dots A_n\}$ is substantively irrational.⁷

What does it mean for it to be “guaranteed” that at least one of $\{A_1 \dots A_n\}$ is substantively irrational? As I understand it, the idea is this: that for each of the possible ways that an agent’s reasons could be, ranging across all (metaphysically) possible worlds and all possible agents, at least one of $\{A_1 \dots A_n\}$ would be substantively irrational. Thus, the hypothesis says that, for any incoherent set of attitudinal mental states of the sort associated with structural rationality, there is no possible world such that *all* of these states are substantively rational. Thus, holding these states will always involve one in substantive irrationality.

An illustrative example of this hypothesis is provided by contradictory beliefs. It’s fairly plausible that, whenever one believes p and believes not- p , at least *one* of these beliefs must be insufficiently supported by one’s (evidence-relative) reasons. For all possible ways that one’s reasons could be, none involve one’s reasons sufficiently supporting both beliefs. The Guarantee Hypothesis tries to extend this thought to all structurally irrational combinations of attitudinal mental states.

If the Guarantee Hypothesis is correct, then a major upshot of this is that (would-be) requirements of structural rationality are normatively *superfluous*, in the sense that they do not add anything to the demands of structural rationality. Any total state that (so-called) structural rationality forbids is already forbidden by substantive rationality. For example, since for any contradictory pair of beliefs, at least one of them is already forbidden by substantive rationality, any further structural requirement forbidding one from believing contradictions would be superfluous. This, the proponents of the Guarantee Hypothesis think, means that we can eliminate these requirements of structural rationality as *sui generis*, distinct requirements. We can account for the intuition that it’s irrational to be incoherent while positing only requirements of substantive rationality.

But suppose that there are conflicts between substantive and structural rationality, where each attitudinal mental state in some set $\{A_1 \dots A_n\}$ is substantively required, but the set as a whole is structurally prohibited. In that case, it immediately follows that the Guarantee Hypothesis is false, since the set is structurally irrational without any of its members being substantively irrational. (Call such cases, then, *conflict counterexamples* to the Guarantee Hypothesis). Thus, conflicts between structural and substantive rationality would close off one of the most prominent and plausible routes to monism.

All of this said, it’s worth me being up front that my ultimate interest is in defending dualism (and rejecting monism), not in establishing the existence of conflicts *per se*. And importantly, there are at least two ways in which putative conflict cases might turn out to threaten monism even if they aren’t ultimately judged to be conflict cases.

First, suppose that it turns out that the right diagnosis of the case is that at least some of $\{A_1 \dots A_n\}$ are not substantively *required*, but merely substantively *permitted*. Nevertheless, if every attitudinal mental state in the set is substantively permitted, but the set as a whole is structurally

⁶ I’ll use this term throughout to refer to both attitudes and absences thereof.

⁷ See Kolodny (2007, 2008a, 2008b), Kiesewetter (2017), and Lord (2018).

prohibited, the Guarantee Hypothesis fails, and the route to monism is again blocked. Call this a *permissive counterexample* to the Guarantee Hypothesis.

Second, suppose that it turns out that the right diagnosis of the case is that substantive rationality actually *forbids* at least one of the attitudes in the set – but the explanation of why this is so invokes facts about the coherence, or *structural* irrationality, of the set. For example, suppose the explanation is that there is a fundamental constraint on reasons that says that, if some combination of responses violate a prior, *sui generis* requirement of structural rationality, then our reasons can't demand that we have all of these responses. If this were so, then the case would not be a counterexample to the Guarantee Hypothesis, but it would show that the Guarantee Hypothesis can no longer be used to plausibly eliminate distinctive requirements of structural rationality. For those very requirements would then turn out to be explanatorily fundamental in explaining why the Guarantee Hypothesis holds, and why conflict and permissive counterexamples to it can't come about. Call any case like this a *fundamentality objection case*.

These other possibilities are important. For I'm going to argue soon that the most plausible ways of resisting the claim that putative conflict cases really are conflict cases end up falling into one of the above categories – thus still undermining monism, or at least versions of it that go via the Guarantee Hypothesis.⁸

3. “Epistemic dilemmas”

One final point before I get into the core of the argument. The volume of which this paper is a part is called *Epistemic Dilemmas*. Are evidence-coherence conflicts epistemic dilemmas? In a loose sense, they obviously are. But I am going to stipulate some terminology to be a bit more precise. Let's say that one is in a *conflict* when one is under two competing requirements that are impossible to jointly satisfy.⁹ And let's say that one is in a *dilemma* when there's no answer to be had as to which of the two competing requirements one (all-things-considered) ought to prioritize satisfying in such a conflict case. It's an open question whether evidence-coherence conflicts are dilemmas in this strong sense. I'll come back to this question in section 6.

Moreover, even if they are dilemmas, there is a sense in which describing them as *epistemic* dilemmas may be misleading. For, at least in my opinion, the requirements of structural rationality on doxastic states aren't *epistemic* requirements in the strictest sense of the term.¹⁰ Let me explain.

In one sense of the term ‘epistemic’ – which is fairly common, but not very faithful to the etymology of the term – an epistemic requirement is just any requirement on *doxastic* states such as

⁸ Notice that I've said in this section only that conflict cases strengthen the case for dualism by undermining the particular version of monism that goes via the guarantee hypothesis. There is a different kind of monism, advanced by Maria Lasonen-Aarnio (2020, 2021) that (perhaps surprisingly) actually allows for evidence-coherence conflicts. I argue against Lasonen-Aarnio's version of monism elsewhere (Worsnip 2021: §3.7).

⁹ This usage is not so weak that it'll describe any kind of opposing *pro tanto* reasons as a conflict. For example, suppose I have a weak reason to jump from a 50-story building, which is that it will give me a momentary thrill, and a very strong reason not to jump, which is that it will kill me. I take it this is not a conflict in any interesting sense. To have a conflict, the two competing normative considerations must amount to *requirements*. Though jumping from the building would give me a thrill, there's no requirement to do so in virtue of this.

¹⁰ Cf. Easwaran & Fitelson (2015) and Smithies (this volume), both of whom describe them as “epistemic requirements”.

belief and credence. Clearly, requirements of structural rationality on doxastic states are epistemic requirements in this loose sense.¹¹

However, in a stricter sense, epistemic requirements are ones that in some way connect with *truth* or *knowledge*, and have to be justified in those terms. For example, on one picture, it has to be shown that obeying these requirements will somehow conduce to true belief or to knowledge. But this stacks the deck against those who think that the structural requirements are just *sui generis* requirements to have attitudes that make coherent internal sense, not to be justified in terms of some further end, either epistemic or practical. For this reason, I prefer not to call them epistemic requirements, and so in a pedantic sense the conflicts I'm describing aren't distinctively *epistemic* conflicts (or dilemmas). It may be that the only specifically *epistemic* rational requirements, strictly speaking, are the substantive ones that require responsiveness to one's evidence. But, I contend, these are not the only rational requirements *on doxastic states*.

4. Preface cases

What I will call "preface cases" are cases based on the structure of the preface paradox.¹² The essential structure of these cases are as follows. You believe a large set of propositions $\{p_1, p_2, \dots, p_n\}$. You have strong evidence for each of these propositions taken individually – enough to make the probability of each proposition, given your evidence, very high. However, for each proposition, this evidence still leaves open a small evidential probability that the proposition is false. Now consider the conjunctive proposition $(p_1 \& p_2 \& \dots \& p_n)$. The risks of error for each individual proposition aggregate for the conjunctive proposition, such that the evidential probability of the conjunctive proposition is very low (given a suitably large n , and given at least some probabilistic independence between the individual propositions).

What should we say about such a case? Let's start with *substantive* rationality. I take the following verdicts to be very hard to resist:

- (1) For each individual proposition in the set $\{p_1, p_2, \dots, p_n\}$, believing this proposition is substantively permissible.
- (2) Believing the conjunctive proposition $(p_1 \& p_2 \& \dots \& p_n)$ is substantively impermissible.

The propositions in the set $\{p_1, p_2, \dots, p_n\}$ are all ones that are highly probable on one's evidence, but that fall short of having probability 1 on one's evidence. (1) follows straightaway if there is some threshold of evidential probability, short of 1, that suffices for rational belief. But, of course, it's

¹¹ Though, it's worth noting that at least in my view, the requirements of structural rationality on doxastic states are very much of a piece with the requirements of structural rationality on practical or conative states, such as the requirement not to have cyclical preferences or inconsistent intentions. The kind of failing involved in violating these various requirement is of exactly the same kind: it involves having mental states that don't fit together with one another in the right ways. It's also notable that there are many requirements of structural rationality that prohibit combinations of doxastic *and* conative states – for example, the requirement that prohibits {believing you ought not Φ , intending to Φ } – and so wouldn't neatly count either as epistemic or practical requirements on this way of talking.

¹² The preface paradox was introduced by Makinson (1965). Important contributions to the contemporary debate about it include, *inter alia*, Ryan (1991), Kaplan (1996: ch. 4), and Christensen (2004: chs. 3-4).

controversial whether this is so. Maybe there are other conditions on rational belief that a proposition with very high evidential probability doesn't always meet.¹³ No matter, though. As long as there are *some* propositions¹⁴ that don't enjoy evidential probability 1 but that it's substantively rational to believe, we can just stipulate, for the purposes of setting up the case, $\{p_1, p_2, \dots p_n\}$ are all propositions of this kind. The claim that there are at least some propositions that don't enjoy evidential probability 1 but that it's substantively rational to believe follows from a very modest kind of fallibilism, according to which the absolute certainty of p on one's evidence isn't a necessary condition for rational belief in p .¹⁵ I think this very modest fallibilism is overwhelmingly plausible, and I will take it for granted here. Thus, (1) can be assumed too.

(2) follows from the assumption that it can't be substantively permissible to believe something that is very improbable on one's evidence. This is overwhelmingly plausible,¹⁶ and I will take it for granted here.

I'll come back to some bolder claims about substantive rationality in a moment. First, though, what about *structural* rationality? There are two putative requirements of structural rationality that might be relevant here:

Deductive Closure Constraint. Where some set of beliefs S entail some proposition q , structural rationality prohibits $\{\text{having each belief in } S, \text{ not believing } q\}$.

Consistency Constraint. Structural rationality prohibits having a set of beliefs S that is inconsistent.¹⁷

As a matter of fact, I think that preface cases put considerable pressure on us to deny the Deductive Closure Constraint *even as a requirement of purely structural rationality*. I've argued more for this elsewhere,¹⁸ but briefly, the thought is this. Suppose you believe each of $\{p_1, p_2, \dots p_n\}$, without having credence 1 in them. Assuming structural rationality requires at least approximate probabilistic coherence, you will be structurally irrational if you don't assign the conjunction $(p_1 \ \& \ p_2 \ \& \ \dots \ p_n)$ a low credence (assuming that $\{p_1, p_2, \dots p_n\}$ are at least fairly probabilistically independent). But, if the Deductive Closure Constraint is true, you will be structurally irrational unless you believe $(p_1 \ \& \ p_2 \ \& \ \dots \ p_n)$. Thus, it seems that if the Deductive Closure Constraint is true, structural rationality requires you to believe

¹³ For example, maybe one's evidence for it has to be "individualized" or not "merely statistical" in some sense (cf. e.g. Thomson 1986; Buchak 2014), that it meets a safety condition (cf. e.g. Enoch, Spectre & Fisher 2012), or that the stakes with respect to it have to be adequately low (cf. e.g. Fantl & McGrath 2002).

¹⁴ Or, more precisely, a non-trivial number of suitably independent propositions.

¹⁵ Timothy Williamson is often understood as denying this, holding that it is, in fact, substantively impermissible to believe something that lacks evidential probability 1. While this may be true of the recent work of his, e.g. Williamson forthcoming, it is actually *not* interpretatively correct as regards his classic *Knowledge and its Limits*, where he explicitly affirms that one can justifiably believe falsehoods, which cannot have evidential probability 1; see Williamson 2000: 9. For a systematic defense of the view that one can justifiably believe that which is not certain on one's evidence, see Brown (2018).

¹⁶ *Pace* Kaplan (1996), and perhaps (more recently) those who defend the thesis that "belief is weak" (e.g. Hawthorne, Rothschild & Spectre 2016). I have argued against Kaplan's view elsewhere (Worsnip 2016).

¹⁷ Either constraint can be weakened to refer only to *recognized* entailments or inconsistencies. I will ignore that complication here, since nothing turns on it for my purposes.

¹⁸ Worsnip 2016, 2021: §9.4.

something for which you have (if you are otherwise structurally rational) a low credence (or to give up your initial beliefs $\{p_1, p_2, \dots p_n\}$).¹⁹ But this seems wrong; on the contrary, it actually seems structurally *irrational* to believe something for which one has a low (<0.5) credence. For given that believing something involving judging it to be true, it is incoherent to believe something while judging it to be more likely to be false than true. Thus, we should reject the Deductive Closure Constraint.

Note that this argument does not appeal to any requirements of *substantive* rationality, or tacitly assume the impossibility of conflicts between the substantive and the structural. It works purely internally to structural rationality, appealing to *other* structural constraints such as probabilistic coherence and the constraint that you not believe something for which you have low (<0.5) credence, and suggesting that the Deductive Closure Constraint would effectively require you to violate one of these constraints.

However, this argument doesn't extend to call the Consistency Constraint into doubt.²⁰ The fundamental reason for this is that, whereas it is structurally irrational to believe something that you have a low credence for (specifically, any credence < 0.5), it isn't structurally irrational to fail to believe something that you have a very high (but non-1) credence for. Thus, no matter how high your credence in the negation of the conjunction gets, failing to believe the negation of the conjunction won't result in any structural irrationality (though it may, if your credence is justified, perhaps be substantively irrational). Thus, I don't think that preface cases show that it's not *structurally* irrational to have inconsistent beliefs. The Consistency Constraint remains plausible.²¹ Or, to put the point more neutrally in the dialectic with the monist who argues for her view via the Guarantee Hypothesis: inconsistency is the sort of pattern *associated* with structural irrationality that the Guarantee Hypothesis should want to account for the irrationality of.

Let's assume, then, that the Consistency Constraint holds. What we would need to generate a conflict counterexample to the Guarantee Hypothesis, then, are the following two claims, bolder than (1) and (2) above:²²

- (3) For each individual proposition in the set $\{p_1, p_2, \dots p_n\}$, believing this proposition is substantively required.
- (4) Believing the *negation* of the conjunctive proposition ($p_1 \ \& \ p_2 \ \& \ \dots \ p_n$) is substantively required.

If (3) and (4), are true, then we're required to believe each proposition in the set $\{p_1, p_2, \dots p_n\}$, but also to believe the negation of the conjunction of these propositions. Since the total set containing $\{p_1, p_2, \dots p_n\}$ plus the negation of their conjunction is inconsistent, this would generate a conflict

¹⁹ This last parenthetical is needed on the assumption that requirements of structural rationality are "wide-scope" (cf. e.g., Broome 2004). I take it that the parenthetical this doesn't make things much better. $\{p_1, p_2, \dots p_n\}$ were just arbitrary beliefs that you have but lack credence in 1. So saying that one is required to give them up would amount to saying that one ought not to have any non-trivial number of beliefs that one lacks credence 1 in. Thus, saying that one is required *either* to believe the conjunctive proposition or to give up one's initial beliefs amounts to saying that one is required *either* to believe something for which one has a very low credence *or* not to have any non-trivial number of beliefs that one lacks credence 1 in. This is still a very bad result.

²⁰ I failed to advertise this fact in Worsnip (2016).

²¹ Or, again, at least a version of it that refers to sets of beliefs one *recognizes* to be inconsistent does.

²² Cf., MacFarlane (ms.), who explicitly takes the view that preface cases are conflict cases.

between substantive rationality and structural rationality (assuming that the Consistency Constraint is a requirement of structural rationality).

If there is some fixed (non-1) threshold of evidential probability, t , above which one is always required to believe a proposition, then it is impossible to resist (3) and (4) – at least given a suitably chosen set $\{p_1, p_2, \dots, p_n\}$, where each proposition's evidential probability exceeds t , but n is large enough such that the probability of the negation of the conjunction also exceeds t . However, this view about when we're required to believe things is, obviously, controversial. It might, then, be thought that it is easy to save the Guarantee Hypothesis from preface-based counterexamples: we just need to deny this threshold-based view about when we're rationally required to believe things.

However, let's consider what the alternatives are. Even if we reject (3), I have already argued that we should accept (1), namely the claim that for each proposition $\{p_1, p_2, \dots, p_n\}$, believing the proposition is substantively permitted. What about if we reject (4)? Suppose we reject (4) but still accept:

- (5) Believing the negation of the conjunctive proposition ($p_1 \& p_2 \& \dots p_n$) is substantively permitted.²³

Given (1) and (5), the preface case turns out to be a *permissive* counterexample to the Guarantee Hypothesis. For it would still be a case in which it's substantively *permissible* to have every belief in the total set containing $\{p_1, p_2, \dots, p_n\}$ plus the negation of their conjunction. Nevertheless, this total set would result in one's violating the Consistency Constraint. So the case would still show that one's violating the Consistency Constraint does not guarantee that one has thereby violated a requirement of substantive rationality. And with this, it would be a counterexample to the Guarantee Hypothesis, even though it's not a *conflict* counterexample to it.

The only other option is:

- (6) Believing the negation of the conjunctive proposition ($p_1 \& p_2 \& \dots p_n$) is substantively prohibited.²⁴

If this is right, then the preface case is neither a conflict counterexample nor a permissive counterexample to the Guarantee Hypothesis.

What could motivate (6), though? That is: *why* does substantive irrationality forbid you from believing the negation of the conjunction, even though its probability on your evidence is very high? And here I think the only forthcoming answer is: because doing so would lead you into irrationality!²⁵

²³ One, but not the only, way to motivate this view would be to hold that it's *never* rationally required to believe a proposition that has a probability short of 1. On this view, cases of high but not-1 probability are permissive cases, where one may permissibly believe or permissibly suspend judgment. Cf. Nelson (2010).

²⁴ Combined with (2), this entails that the only permissible attitude to take toward the conjunctive proposition (and its negation) is to suspend judgment on it.

²⁵ An anonymous referee suggested that it might be because the evidence for the negation of the conjunctive proposition is merely statistical, and belief on merely statistical evidence is (on some accounts; cf. fn. 13) impermissible. I don't think this is so (at least not in any sense on which belief on statistical evidence might be impermissible). The evidence for the negation of the conjunctive proposition just is body of evidence that one has for the conjuncts; while this body of

But if this is so, then structural rationality is effectively putting a constraint on substantive rationality: because inconsistent beliefs are irrational, our theory of substantive rationality is being constrained so as not to permit all of them. Since $\{p_1, p_2, \dots p_n\}$ are just arbitrary propositions that we can stipulate that it's substantively rational to believe, the only way to do this is to forbid belief in the negation of the conjunction.

Thus, on this diagnosis, the preface case turns out to be what I earlier called a fundamentality objection case. In such a case, the Guarantee Hypothesis is saved from counterexample only by making a requirement of structural rationality (viz. the Consistency Constraint) fundamental to the explanation of why it holds, by putting a constraint on substantive rationality. This undermines the attempt to eliminate structural rationality or its requirements, as the monist wants to do.

I have no settled view on which of the above views is the right one (beyond insisting on the truth of (1) and (2)). But I have argued that for all the logically possible ways of diagnosing the case, it turns out either to be a conflict counterexample, a permissive counterexample, or a fundamentality objection case. All three of these possibilities foreclose the possibility of monism. If my argument is sound, the only way for the monist to reply would be to simply deny that inconsistent beliefs are irrational in *any* sense of the term 'irrational', and to exclude them from the set of "attitudes associated with structural irrationality" that the Guarantee Hypothesis was supposed to account for. But this should at the very least be a disappointment for the monist, who would – I take it – have hoped to include inconsistent beliefs among the patterns of irrational attitudes that she can account for by appeal to the Guarantee Hypothesis. Strict logical inconsistency of belief is, after all, perhaps the paradigm case of incoherence.

Of could, the monist could just insist that the preface case itself undermines the Consistency Constraint. But I think this begs the question. Let me explain.

Suppose we didn't have the possibility of dualism about rationality in view. We might then – as many epistemologists have done – assume that having doxastic attitudes that are supported by the evidence on its own is not just necessary but also *sufficient* for those doxastic attitudes being rational (since we would be ignoring the possibility of *sui generis* coherence requirements that go beyond the requirement to believe what one's evidence supports). Turning to the preface case, then, we might think that if one's evidence at least adequately supports each of the relevant jointly inconsistent beliefs, then it straightaway follows that inconsistent beliefs can be fully rational, and that there's no kind of consistency requirement on belief.²⁶

But once the dualist view is on the table as an option, this is too quick. The dualist posits separate coherence (structural) requirements on belief that may add to the demands of substantive rationality, that is, of evidence-responsiveness. So if dualism is a live possibility, we cannot conclude simply from the fact that each of the inconsistent beliefs is supported by the evidence that it is not irrational in *any* sense (even a structural one) to have inconsistent beliefs. It may be that the

evidence makes all each conjunct highly probable, it also (by failing to make the conjuncts certain) makes the negation of the conjunction very probable. (True, one might have to do some statistical reasoning to *see* that it makes the negation of the conjunction very probable. But that does not make the evidence itself statistical in the relevant sense, I think.) Thus, if the evidence isn't merely statistical in a way that would make it impermissible to believe the conjuncts, it's always not merely statistical in a way that would make it impermissible to believe the negation of the conjuncts.

²⁶ Cf., e.g., Christensen (2004).

inconsistent beliefs are substantively permissible, but not structurally permissible. (This need not, it bears stressing, automatically amount to positing a conflict between substantive and structural rationality. As we've seen, this follows only given the further, bolder claim that the inconsistent beliefs are substantively *required*.)

So we now see that the argument against the Consistency Constraint from preface cases turns on the monist assumption that the evidential requirements on belief are the only rational requirements on belief, and that there are no *sui generis* structural requirements.²⁷ Thus, to insist on it in defense of monism against the challenge I've made to it would simply beg the question.

5. Misleading higher-order evidence cases

Let's now turn to cases of misleading higher-order evidence. Higher-order evidence, roughly, is evidence about your evidence (either about what your evidence is, or what it supports).²⁸ But higher-order evidence can be misleading. Consider, for example, the following case (from Worsnip 2018):²⁹

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 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."

Here's what should be fairly uncontroversial in this case. First, taken on their own the clues do not support believing (the first-order proposition) that the vicar did it. Second, taken on its own Miss Marple's testimony supports believing (the higher-order proposition) that the clues support believing that the vicar did it. In that sense, the case involves misleading higher-order evidence.

But the controversial question, for the debate about higher-order evidence, is what to say about what the *total* evidence supports in this case. Let's say that the case is a case of *all-things-considered* misleading higher-order evidence if one's total evidence (clues + testimony) supports believing a falsehood about what one's total evidence (clues + testimony) supports.

²⁷ In this way, it's *unlike* the argument against the Deductive Closure Constraint from preface cases that I myself endorsed above. As I was at pains to point out, that argument works purely internally to structural rationality, and makes no tacit monist assumptions. But it doesn't extend to the Consistency Constraint.

²⁸ For a more precise characterization, see Chen & Worsnip (forthcoming).

²⁹ See also Coates (2012) and Horowitz (2014) for similar cases.

Such cases, if they exist, can generate conflicts between substantive and structural rationality. Consider again the Miss Marple & Mabel case. Suppose that Mabel's *total* evidence (decisively³⁰) supports suspending judgment about whether the vicar did it. Then:

(7) Mabel is substantively prohibited from believing that the vicar did it.

But suppose that Mabel's total evidence (decisively) supports believing that her total evidence (decisively) supports believing that the vicar did it. Then:

(8) Mabel is substantively required to believe that her total evidence (decisively) supports believing that the vicar did it.

However, the following is very plausibly a requirement of structural rationality:

Inter-Level Coherence (Pos).³¹ Structural rationality prohibits {believing that your total evidence decisively supports believing p , not believing p }.³²

If (7) and (8) are true, then substantive rationality requires Mabel to violate Inter-Level Coherence, and so we have a conflict counterexample to the Guarantee Hypothesis.

Now, it is difficult to conclusively establish that the Miss Marple and Mabel case, as described, is a case of all-things-considered misleading higher-order evidence, and hence that in this case as it is described. After all, Miss Marple's testimony is not *just* evidence for the higher-order proposition that the evidence supports believing that the vicar did it. It is also at least some evidence for the first-order proposition that the vicar did it. When Miss Marple says that the evidence supports p , that's some indication not just that the evidence supports p , but also that p . (Moreover, slightly more arguably, the clues themselves may constitute some evidence against the higher-order proposition that the evidence supports believing that the vicar did it.) Whether the case constitutes an example of all-things-considered misleading higher-order evidence thus depends on the exact relative weights of the clues and the testimony, first, at the level of the first-order proposition that the vicar did it, and second, at the level of the higher-order proposition that the evidence supports believing that the vicar did it. And those exact relative weights are hard to specify.

However, in my (2018), I offered the following argument that at least *some* cases of all-things-considered misleading higher-order evidence should be possible. Though Miss Marple's testimony has some weight at the level of both the higher-order proposition and the first-order proposition, it should

³⁰ By 'decisively' supporting p , I mean that it supports it in such a way as to make believing p substantively required, not just substantively permissible. We could talk of 'adequately' supporting p to refer to support that makes believing p substantively permissible. Of course, on non-permissivist theories, decisively and adequate support are one and the same thing, and we can omit these qualifiers.

³¹ I call it 'pos' to distinguish it from the (equally plausible) negative inter-level coherence requirement that prohibits {believing that your total evidence doesn't adequately support believing p , believing p }. This principle can generate substantive-structural conflicts in roughly the same way, but it's the positive principle we need for the Miss Marple and Mabel case specifically.

³² I argue for this principle in Worsnip (2018).

normally have *more* weight at the level of the higher-order proposition than at the level of the first-order proposition. After all, what Miss Marple says is that the evidence supports believing that the vicar did it. Intuitively, this bears *more directly* on the proposition that the evidence supports believing that the vicar did it than it does on the proposition that the vicar did it. So Miss Marple's testimony is *more weighty* at the level of the higher-order proposition than it is at the level of the first-order proposition.

If you doubt that testimony of the sort that Miss Marple gives (or any evidence at all) could be more weighty at the level of the higher-order proposition than it is at the level of the first-order proposition, consider the following related case. Suppose there is an epistemic oracle that you know to be infallible. But the epistemic oracle only ever tells you about higher-order facts about what your total evidence supports; it never just directly tells you the first-order facts. Suppose the epistemic oracle tells you that your evidence supports believing that the vicar did it. If the epistemic oracle really is infallible, and you can be certain of this, then the probability on your evidence that the evidence supports believing that the vicar did it is 1. But is the probability on your evidence *that the vicar did it* 1? No. For the oracle's infallibility consists only in its always delivering truths about what the evidence supports. And evidence is sometimes misleading.³³ So even given the oracle's infallibility, and the fact that the oracle says that your evidence supports believing that the vicar did it, there is still some non-zero probability that the vicar didn't in fact do it.

Now, the oracle case is not itself a case of misleading total evidence about what your total evidence supports.³⁴ But it does illustrate how testimony about higher-order matters, and higher-order evidence more generally, can probabilify a higher-order proposition more strongly than it probabilifies a first-order proposition. Given the oracle's testimony, the probability of the higher-order proposition that the evidence supports believing that the vicar did it is 1, but the probability of the first-order proposition that the vicar did it falls short of 1. And this result continues to hold when we relax the assumption of infallibility that's built into the oracle case, and switch to a case where the testifier is just extremely reliable. Elementary calculations show that, if we set up the prior probabilities in the right way, Miss Marple's testimony can also probabilify the higher-order proposition that the evidence supports believing that the vicar did it to a higher degree than it probabilifies the first-order proposition that the vicar did it (see Worsnip 2019 for the details).

Suppose, then, that Miss Marple's testimony is more weighty at the level of the higher-order proposition than the first-order proposition. This suggests that it should at least *in principle* be possible for Miss Marple's testimony to be *decisive* with respect to the higher-order proposition (that the evidence supports believing that the vicar did it) – that is, to make it probable enough to be worthy of belief – without being decisive with respect to the first-order proposition (that the vicar did it). It won't just automatically be true that whenever it's weighty enough to be decisive with respect to the higher-order proposition, it's weighty enough to be decisive with respect to the first-order proposition.

³³ I'm not begging the question here. What's at issue is whether one's total evidence can mislead about what one's total evidence supports. But it's common ground that you can sometimes get misleading evidence about ordinary first-order propositions like whether the vicar did it.

³⁴ The reason is this. In this case, when the oracle says that that the evidence supports believing that the vicar did it, this proposition must be true, by the setup of the case. Unlike Miss Marple, the oracle cannot be making an uncharacteristic mistake.

Of course, whether Miss Marple's testimony actually *is* decisive with respect to the higher-order proposition, but not the first-order proposition, in any particular concrete instance of the case depends on how the case is filled in. In particular, it depends on how much evidence is needed for rational belief (in the relevant context), as well as how evidentially probable the relevant propositions are before Miss Marple's testimony. But if Miss Marple's testimony is more weighty with respect to the higher-order proposition than it is with respect to the first-order proposition, we should expect there to be *some* version of the case where it is decisive with respect to the former but not with respect to the latter. Thus we should expect there to be at least some precisification of the Miss Marple and Mabel case that is a case of all-things-considered misleading higher-order evidence, where (7) and (8) are true.³⁵ If so, these will be conflict counterexamples to the Guarantee Hypothesis.

Now, just as in preface cases, one might push back against this diagnosis. The foregoing tacitly assumed that, if the higher-order proposition that the evidence supports believing that the vicar did it is more probable on the evidence than the first-order proposition that the vicar did it, then at least in principle, one might be substantively required to believe the former proposition while being substantively required not to believe the latter. However, again, if it's not the case that there's some fixed threshold of evidential probability above which one is substantively required to believe a proposition, then there is some room to question this.³⁶

As in the preface case, however, I now want to argue that one can save the case from counting as a conflict counterexample only at the expense of turning it into a permissive counterexample or a fundamentality objection case. Suppose first that one holds that even though the higher-order proposition is very probable, Mabel is not substantively *required* to believe it, but merely permitted to do so. Then we replace (8) with:

- (9) Mabel is substantively permitted to believe that her total evidence (decisively) supports believing that the vicar did it.

But then the case is a permissive counterexample to the guarantee hypothesis. For it would still be substantively *permissible* for Mabel to {believe that her total evidence (decisively) supports believing that the vicar did it, not believe that the vicar did it}, and thus to be structurally irrational. The problem likewise remains if we also say that believing that the vicar did it is not forbidden, but merely *not required*, such that we replace (7) with:

- (10) Mabel is substantively permitted to believe that the vicar did it, and substantively permitted not to believe that the vicar did it.

For even if this is so, it would still be substantively permissible for Mabel to {believe that her total evidence (decisively) supports believing that the vicar did it, not believe that the vicar did it}.

³⁵ For another defense of the possibility of all-things-considered misleading higher-order evidence, see Lasonen-Aarnio (2020).

³⁶ See Lord (2018: 60) for this response to my case. Ram Neta also suggested the same thing to me independently.

The only way to avoid diagnosing the case as a counterexample to the guarantee hypothesis would be either to claim that Mabel's potential higher-order belief (viz. that the total evidence supports believing that the vicar did it) and her potential first-order belief (viz. that the vicar did it) are either *both* substantively required, or that they are *both* substantively prohibited. But, again, what would *explain* why this has to be so, given that the higher-order proposition is more probable on Mabel's evidence than the first-order proposition is? Again, I think the only forthcoming answer seems to be: because structural rationality puts some kind of constraint on what substantive rationality permits.³⁷ So, for example, perhaps the reason why it can't be substantively rational for Mabel to suspend judgment on whether the vicar did it is precisely *because* doing so would make her inter-level incoherent, and thus structurally irrational. If we invoke this explanation, then again, the case becomes a fundamentality objection case. But without such an explanation, it looks like a miraculous or magical coincidence that even though the higher-order proposition is more probable than the first-order proposition on Mabel's evidence, their normative status (as substantively required or prohibited) *must* travel along together in every possible precisification of the case.

One popular strategy for trying to explain why the normative status of higher-order and first-order beliefs must travel together – which avoids appealing to any fundamental claims about structural rationality – is to appeal to some view about “higher-order defeat” according to which, if one has a justified higher-order belief that one's first-order belief is unjustified, then this defeats the justification one has for one's first-order belief, or provides reason to suspend judgment at the first-order level.³⁸ This may work for cases where one has first-order evidence that supports believing *p*, but has a justified (but false) higher-order belief that one's first-order evidence *doesn't* support believing *p*. But it doesn't generalize to cover the Miss Marple and Mabel case, where things are the other way round: the first-order evidence *doesn't* support believing *p* (i.e., that the vicar did it) in the first instance, and Mabel's justified (but false) higher-order belief is that this evidence *does* support believing *p*. There's no justification for (first-order) believing for this justified higher-order belief to defeat.³⁹ Similarly, what's needed to line up the justificatory status of one's higher-order and first-order beliefs are reasons to *believe*, not reasons to suspend (at the first-order level). So this explanation of the case won't work.

Just like in the preface case, then, the diagnosis of Mabel's case as a conflict counterexample can only be avoided by turning it into a permissive counterexample or a fundamentality objection case. One way or another, again, the case spells trouble for monism.

Once again, the defender of monism could claim that cases of misleading higher-order evidence provide us with reason to simply reject Inter-Level Coherence as a requirement of rationality,

³⁷ Cf. Neta (2018: 327).

³⁸ Bergmann 2005, Skipper 2019, González de Prado 2020, Lord & Sylvan 2021.

³⁹ Could we say that the justified higher-order belief defeats Mabel's justification for *suspending judgment* that the vicar did it? In order for this to prevent the case from being a conflict or permissive case, it would have to make it the case that Mabel now *is* justified in believing that the vicar did it. But it's dubious that the state of suspending judgment about *p* can be defeated, and *belief in p thereby justified*, except by evidence for *p*. This represents an asymmetry with the reverse situation: belief in *p* can be defeated, and suspended judgment made rational, by considerations that are not evidence against *p*. I think this is because believing is not a default state in the way that suspending judgment is: whereas a mere absence of reasons to believe make suspending judgment rational, a mere absence of reasons to suspend cannot make belief rational; only adequate evidence *for p* can do that. See Tucker (forthcoming).

structural or otherwise.⁴⁰ But as in the preface case, this begs the question. If believing what one's evidence supports *sufficed* for rationality, as the monist holds, then it would follow from the existence of cases where one's evidence supports a falsehood about what one's evidence supports that one can rationally violate Inter-Level Coherence. But with dualism on the table as a live option, this inference doesn't go through. It may be that one can violate Inter-Level Coherence without *substantive* irrationality, but that it is still structurally irrational to do so.

There's one other way for the monist to try to resist the challenge from higher-order evidence cases that wasn't available in preface cases. I've presumed, in presenting the Miss Marple and Mabel case, that Miss Marple's testimony that the clues support believing the vicar did it has evidential weight, and that they could in principle make it the case that substantive rationality requires or at least permits Mabel to believe that her total evidence supports believing that the vicar did it. But a very extreme *a priorist* view denies this.⁴¹ This view holds that evidential support relations are *a priori*, and moreover (and more strongly) that in virtue of this, Mabel possesses *indefeasible* *a priori* justification for true beliefs (and against false beliefs) about what her evidence supports, that cannot be defeated by any misleading empirical evidence (e.g. testimony) to the contrary. If this is so, then as long as Mabel's total evidence doesn't support believing that the vicar did it, she cannot be substantively required or even permitted to believe that her total evidence doesn't support believing that the vicar did it.

Even granting that evidential support relations are wholly *a priori*, I find the view that we enjoy indefeasible justification for *a priori* questions, that can never be empirically defeated, wild. On this view, I should be completely certain in all the epistemological, moral, and more generally *a priori* philosophical truths – some of the hardest and most disputed questions there are in the canons of human thought. However, even if this view were plausible, there is a further issue for this strategy. Cases of all-things-considered misleading higher-order evidence can be generated in two ways. One is via misleading evidence about the evidential support relations – that is, how strongly the items in one's evidence support different potential doxastic attitudes. The second is via misleading evidence about what one's evidence *is* in the first place – that is, which propositions (or other items) are part of one's evidence in the first place.⁴² But even if the evidential support relations are wholly *a priori*, it's not (wholly) *a priori* which items are part of one's evidence.⁴³ Thus, the strong *a priorist* story can't generalize to block all potential cases of all-things-considered misleading higher-order evidence.

6. How to respond to evidence-coherence conflicts: misleading higher-order evidence cases

Let's now suppose that preface and/or misleading higher-order evidence cases *are* best-understood as conflicts of evidence and coherence. What then?

⁴⁰ Cf. Coates (2012); Lasonen-Aarnio (2020).

⁴¹ Cf. Titelbaum (2015); Smithies (this volume).

⁴² For some cases of all-things-considered misleading higher-order evidence generated by misleading evidence about what one's evidence *is*, see Worsnip (2018).

⁴³ For example, suppose that, as Williamson (2000) contends, one's evidence is one's knowledge. If this is so, then one knows that some proposition is part of one's evidence by knowing that one knows it. But it's clearly not true that knowledge that one knows *p* is wholly *a priori* for all *p* – if one's knowledge of *p* is *a posteriori*, typically one's knowledge that one knows *p* is going to be at least partly *a posteriori* too.

As I said earlier, evidence-coherence conflicts will only be *dilemmas* in the strict sense I defined if there is no answer as to whether, in such cases, it's the evidential (substantive) requirements or the coherence (structural) requirements take precedence. In this section, I want to argue that, in cases of misleading higher-order evidence, there is such an answer: the structural requirement to avoid inter-level incoherence takes precedence over the substantive requirement to take the doxastic attitudes that one's evidence supports. This is because, as I will argue, in these cases the agent has an *excuse* for violating the substantive requirements, but not for violating the structural one.⁴⁴ I'll offer two (closely related) arguments for this.

6.1 *The argument from unknowability of one's predicament*

Here's a crucial fact about cases of all-things-considered misleading higher-order evidence like Mabel's: you cannot *know* that you are in such a case. For it's part of what it is for evidence to be misleading that you don't know it's misleading.

Let me present this a little more precisely. Recall that, if Mabel's case is a case of all-things-considered misleading higher order evidence, then:

- (i) Mabel's evidence supports believing that her evidence supports believing that the vicar did it; but
- (ii) Mabel's evidence doesn't support believing that the vicar did it.

Suppose for reductio that Mabel knows that she is in such a situation; that is, that she knows (i) and (ii). Then, as well as (i) and (ii) holding, the following also hold:

- (iii) Mabel knows that her evidence supports believing that her evidence supports believing that the vicar did it; and
- (iv) Mabel knows that her evidence doesn't support believing that the vicar did it.

However, (i) and (iv) are not co-possible. This follows from the following general principle:

No Knowledge Against Your Evidence. If your evidence supports believing p , then you can't know not- p .

Thus, it's impossible for Mabel to know that she's in a case of all-things-considered misleading higher-order evidence. A similar argument will show that Mabel can't justifiably believe that she's in a case of all-things-considered misleading higher-order evidence, given the following principle:

No Justified Belief Against Your Evidence. If your evidence supports believing p , then you can't justifiably believe not- p .

⁴⁴ Excuses are distinct from justifications here. (Cf., *inter alia*, Fletcher 1985; Baron 2005; Williamson forthcoming.)

Now, here is what I think is another very plausible principle:

Unknowability Is an Excuse. If an agent could not have possibly known (or justifiably believed) that she was under some normative demand(s), then she has an excuse for failing to (fully) satisfy that/those demand(s).⁴⁵

Given this principle, it follows that since Mabel couldn't have known that substantive rationality requires her both to believe that her evidence supports believing that the vicar did it, and to refrain from believing that the vicar did it, she has an excuse for failing to satisfy (both of) these demands.

By contrast, Mabel has no similar excuse for violating Inter-Level Coherence. This requirement applies to her under all circumstances, so she doesn't need to know anything special about her circumstances to know that it applies to her here. And even if she doesn't have the terminology of "structural irrationality" specifically, there's no in-principle barrier to her knowing that there would be something incoherent and irrational about her believing that the evidence supports believing that the vicar did it, and yet not believing that the vicar did it.⁴⁶

Thus, Mabel has an excuse for violating the substantive requirement, but not the structural requirement. The final principle we now need is the following:

Excuses Resolve Dilemmas. If one is required to Φ , and one is required to Ψ (where these two requirements are not co-satisfiable), but one has an excuse for failing to Φ and no excuse for failing to Ψ , then there is a good sense in which one ought to Ψ rather than Φ .

Given this, it follows that there's a good sense in which Mabel ought to satisfy the requirement of structural rationality, rather than those of substantive rationality, when they conflict.⁴⁷ Depending on how one wants to use the term 'dilemma', this either *dissolves* the dilemma that Mabel faces, or at the very least tells us how she ought to respond to it.

6.2 *The argument from the impossibility of ex post rationality*

Consider now the distinction between *ex ante* and *ex post* (substantive) rationality. On the standard account of this distinction, very roughly speaking, an attitude is *ex ante* substantively rational when your reasons support it (in the case of a doxastic attitude: when your *evidence* supports it), and *ex post* substantively rational when you hold the attitude in a way that constitutes properly responding to those very reasons (evidence) that support it. A belief can be *ex ante* rational without being *ex post* rational, but it can't be *ex post* rational without being *ex ante* rational.

⁴⁵ This is much weaker than saying that mere ignorance of the normative demands one is under is an excuse (cf. e.g. Harman 2011). One can be ignorant of a demand for all sorts of reasons other than the fact that one *couldn't possibly* have known that one was under it.

⁴⁶ Nor can she justifiably believe that this is OK *because she's in a case of misleading higher-order evidence*, since, as we've just seen, she can't know or justifiably believe that she's in a case of misleading higher-order evidence.

⁴⁷ Here I'm assuming that there's a good sense in which one ought to satisfy the requirements of structural rationality in the first place. See Worsnip (2021: ch. 8) for a treatment of this issue.

Now, cases of all-things-considered misleading higher-order evidence are ones where your evidence supports incoherent attitudes. But what follows from this is that in such cases, it's *ex ante* substantively rational to have incoherent attitudes. It does not follow from this that in such a case, you can have these incoherent attitudes and be *ex post* substantively rational in so doing.

In fact, Paul Silva (2017) has compellingly argued that you cannot. Here is an argument inspired by his. Suppose that Mabel believes, *ex ante* and *ex post* rationally, that her evidence (decisively) supports believing that the vicar did it. In that case, she believes that her evidence does not permit suspending judgment about whether the vicar did it. But the following principle is plausible:

Contrary Judgments Preclude Proper Responses. If an agent believes that her reasons do not permit Φ -ing, then even if her reasons *do* permit Φ -ing, she can't Φ in a way that constitutes properly responding to her reasons (so as to be *ex post* substantively rational in Φ -ing).

To see the case for Contrary Judgments Preclude Proper Responses, consider the following. Properly responding to some reason for Φ -ing, in such a way as to be *ex post* rational, doesn't *just* require Φ -ing on the basis of some reason, R, that supports Φ -ing.⁴⁸ Suppose that the fact that you promised to meet me for lunch is a decisive reason to show up. And suppose this is so in virtue of its being morally required for you to keep your promises. But suppose that you instead think it's a decisive reason to show because you, falsely and irrationally, believe that anyone who fails to keep a promise will at some point later in their life be hit by a tornado. If you then go on to show up to lunch, even if you do so in a good sense *on the basis of* your having promised to do so, you don't count as having properly responded to your reason for doing so; nor does your showing up count as *ex post* rational. What is also needed is some sufficient degree of understanding of *why* your promise counts in favor of showing up. More generally, properly responding to a reason R for Φ -ing requires not just Φ -ing on the basis of R but having some degree of understanding of why R supports Φ -ing. But you can't *understand why your reasons support Φ -ing* if you yourself believe they don't even permit Φ -ing. From this, Contrary Judgments Preclude Proper Responses follows.

Apply this to Mabel's case: since Mabel herself believes that suspending judgment is the wrong attitude to have given her evidence, by Contrary Judgments Preclude Proper Responses, she can't be suspending judgment in a way that constitutes properly responding to her substantive reasons to suspend judgment. And so she is not *ex post* (substantively) rational in suspending judgment.⁴⁹ Thus, if Mabel has the higher-order attitude that substantive rationality requires her to have, then she can't be *ex post* rational in holding the first-order attitude that substantive rationality requires her to have.⁵⁰ Thus, she cannot be *ex post* rational in holding both attitudes.

⁴⁸ Cf. Singh (ms.).

⁴⁹ I see no reason to deny that suspended judgment admits of the *ex ante*/*ex post* rationality distinction. As Friedman (2013) has influentially argued, suspended judgment is not merely absence of belief but an attitude in its own right. And even when one's reasons favor suspending judgment, one can either suspend judgment for the reasons that favor it, or for unrelated, bad (motivational) reasons.

⁵⁰ Some may want to put the point by saying that the *ex post* rationality of Mabel's state of suspending judgment about whether the vicar did it is *defeated* by her higher-order belief. I don't object to this way of talking, but it's important not to confuse the contention here with the claim (already discussed and criticized in fn. 39 above) that Mabel's higher-order

Silva himself thinks that this appeal to *ex post* justification “dissolves the puzzle of misleading higher-order evidence”, and that in particular it gets rid of the sense that such cases are dilemmatic in any troubling way. But for all that has been said so far, it’s not clear why this should be so. Indeed, one might think that the above points actually secure a way of presenting the dilemma without even appealing to any requirement of structural rationality. For consider Mabel’s options with respect to the proposition that the vicar did it. By hypothesis, her evidence doesn’t support believing this proposition, or disbelieving it (i.e. believing its negation). Assuming, as is standard, that an attitude is *ex post* rational only if it’s *ex ante* rational, this means that neither of these attitudes are *ex post* substantively rational. But if the argument from the last paragraph is right, then (at least if she has the higher-order attitudes that it’s substantively rational to have) she *also* can’t be *ex post* substantively rational in suspending judgment about this proposition. So it seems like, whatever attitude Mabel takes toward the proposition that the vicar did it, it will fail to be *ex post* substantively rational.

However, again, I think there is a plausible principle about *excuses* that can be invoked here to soothe the dilemma:

Impossibility of Rationally Responding is an Excuse. If there is no possible way for one to comply with an (*ex ante*) demand of (substantive) rationality in such a way as to be *ex post* (substantively) rational, then one has an excuse for failing to satisfy this demand of (substantive) rationality.

If this is right – and again, I think it is very plausible – then Mabel has an excuse for failing to satisfy the substantive requirement to suspend judgment about whether the vicar did it.⁵¹ Again, she has no similar excuse for failing to satisfy the requirement of inter-level coherence. Indeed, it’s not clear that the *ex ante/ex post* distinction makes sense for structural rationality. A response is *ex post* rational, roughly, if constitutes a correct response to the reasons that make the response *ex ante* rational. But since structural rationality isn’t about responding to reasons, it’s unclear how to apply the distinction to structural rationality. So again, by Excuses Resolve Dilemmas, Mabel ought to satisfy the requirement of structural rationality rather than the requirement of substantive rationality here.

Interestingly, this second argument allows us to be a bit more precise about the response that Mabel ought to have in her case. Assuming that the requirement of inter-level coherence is “wide-scope,”⁵² there are *two* ways for Mabel to satisfy it:

- a) {Believe that her evidence supports believing that the vicar did it, believe that the vicar did it};
- b) {Not believe that her evidence supports believing that the vicar did it, not believe that the vicar did it}

evidence defeats her *ex ante* justification for suspending judgment about whether the vicar did it in a way that makes believing that proposition positively substantively rational.

⁵¹ And, if (as is plausible) it’s not an option for Mabel to have *no attitude whatsoever* towards the proposition that the vicar did it – not even suspending judgment – then she likewise has an excuse for failing to satisfy the requirement *not* to believe that the vicar is guilty.

⁵² Again, see e.g. Broome (2004).

The Argument from the Unknowability of One's Predicament seemed to be neutral on these two options. But the Argument from the Impossibility of *Ex Post* Rationality is not. Of the two things that substantive rationality requires of Mabel – believing that her evidence supports believing that the vicar did it, and suspending judgment about whether the vicar did it – it is specifically the latter attitude that Mabel can't have *ex post* rationally in light of the former, not the reverse. This suggests that she has an excuse for failing to suspend judgment about whether the vicar did it, but no excuse for failing to believe that her evidence supports believing that the vicar did it. This suggests that Mabel has an excuse for violating the requirement of substantive rationality in the way involved in option (a), but not for violating it in the way involved in option (b), and thus that she in a good sense ought to adopt the attitudes in option (a).

6.3 *Why only an excuse?*

Someone might now ask: why only say that the foregoing considerations show that one has an *excuse* for violating the requirements of substantive rationality in cases of all-things-considered misleading higher-order evidence? Why not take them to show something stronger, namely that there is no requirement of substantive rationality that one would be violating by having structurally rational attitudes in these cases? This would seem to follow if we replace the principles Unknowability Is an Excuse and/or Impossibility of Rationally Responding Is an Excuse with the following stronger principles:

Luminosity. An agent is under some normative demands(s) only if it is possible for her to know (or justifiably believe) that she is under these demand(s).

Rational Respondability. An agent is *ex ante* substantively required to Φ only if it possible for her to Φ *ex post* rationally.

If these principles are true, then not only would the cases turn out not to be *dilemmas* (in my strong sense) – they wouldn't even turn out to be *conflicts* between substantive and structural rationality at all.⁵³

I am skeptical about Luminosity and Rational Respondability. Addressing them directly, though, is a task beyond the confines of this essay.⁵⁴ So I will make do with pointing out how revisionary the result would be if we apply these principles in cases of all-things-considered misleading higher-order evidence. Specifically, the result would be to say that substantive rationality no longer requires us to take the doxastic attitudes that our evidence supports in such cases. If correct, this would be a result of immense significance. I leave it to philosophers to decide which they would rather hang on to: the claim that substantive rationality requires us to take the doxastic attitudes that our

⁵³ Though they'd still be conflicts between *evidence* and coherence in a good sense. They would just be cases where substantive rationality doesn't require one to have the attitudes supported by one's evidence.

⁵⁴ For skeptical perspectives on Luminosity, see, *inter alia*, Hawthorne & Srinivasan (2013) and Srinivasan (2015).

evidence supports, or principles like Luminosity and Rational Responsibility. If there are cases of all-things-considered misleading higher-order evidence, they cannot have both.

7. Why this doesn't work for preface cases

As we have seen, there is a nice story about why, in conflict cases generated by misleading higher-order evidence, the demands of structural rationality take precedence. It would be nice if this story generalized to preface cases.

Unfortunately, it doesn't seem to do so. This is because the considerations about unknowability and *ex post* rationality that applied in misleading higher-order evidence cases do not apply in preface cases. Suppose that preface cases do constitute conflict cases. Unlike in the case of misleading higher-order evidence, there is nothing to prevent you, in principle, from knowing that you're in a preface case, and that you're under conflicting demands in such a case. You might know that each of the individual propositions $\{p_1, p_2, \dots p_n\}$ is highly probable, and that the negation of their conjunction is highly probable. If we're required to believe propositions that exceed some threshold of probability, you could also in principle know this, and thus know that substantive rationality requires you to believe inconsistent propositions in the preface case. But still, you could also know in principle that it's structurally irrational to do so.

Slightly more arguable, it seems possible to respond to the demands of substantive rationality in an *ex post* rational way in preface cases. You could believe each proposition $\{p_1, p_2, \dots p_n\}$ on the basis of its being highly probable on your evidence, and likewise for the negation of the conjunction. This would seem to suffice for *ex post* substantive rationality in each belief.

Now, someone might claim that if you believe $\{p_1, p_2, \dots p_n\}$, that thereby precludes your belief in the negation of their conjunction from being *ex post* substantively rational. This is an intriguing claim. Effectively, this view says that believing something on the basis of the considerations that make it *ex ante* rational does not suffice for *ex post* rationality. If this can be vindicated, then perhaps we have a way of running a version of The Argument from the Impossibility of Ex Post Rationality to show that the structural requirement against inconsistent beliefs takes place in the preface case. But this is a major task to execute. Thus, it's at least more complicated whether preface-type conflict cases can be handled in the way that I suggested we handle misleading higher-order evidence cases. If they cannot be – and if, as I've been supposing for the sake of argument, preface cases really are best-understood as conflict cases – they will be genuine dilemmas in a stronger sense.

8. Conclusion

Let us take stock. I argued in §§4-5 that preface cases and misleading higher-order evidence cases are both plausibly diagnosed as conflict cases, but that there are also ways to resist these diagnoses. However, I also argued that the most plausible *alternative* diagnoses of these cases still end up putting pressure on monism about rationality, just as strongly as a diagnosis of the cases as conflict cases would. Thus, one conclusion of this chapter is that these cases put strong pressure on monism about rationality – even if the diagnosis of such cases as conflict cases can be resisted.

In §§6-7, I took up the question of whether, assuming these cases *are* best understood as conflict cases, they are *dilemmas* in a stronger sense: where a *conflict* involves being under conflicting requirements, and a *dilemma* involves being in a situation where there's no fact of the matter about which of the thing conflict requirements one ought, all-things-considered, to satisfy. I gave two arguments – the Argument from Unknowability, and the Argument from the Impossibility of Rational Responsibility – for the view that misleading higher-order evidence cases, even if best-understood as conflict cases, are *not* dilemmas in this stronger sense. But I also noted that these arguments don't seem to extend to preface cases. So, *if* preface cases are best understood as conflict cases, it is hard to see how to avoid concluding that they are also dilemmas.

How should we react to this? A crucial point here is that preface cases (as I've defined them) are *not* just *recherché* cases that arise for certain historians writing a certain kind of preface to a certain kind of book. Rather, we are all in cases with preface-like structures right now. All that is required to be in such a case is to have a non-trivial number of beliefs, for propositions for which we have credence at least somewhat short of 1, in propositions that are at least somewhat independent. Once those conditions are met, the conjunction of all these propositions will be extremely improbable, and that raises the question of whether to believe the (extremely probable) negation of this conjunction. On the conflict diagnosis of the case, to believe it would make one structurally rational, but to refrain from believing it would make one substantively irrational.

But one may feel queasy about this diagnosis. Our predicament as believers is that we need to believe a fairly large number of things, about a range of (somewhat independent) subject matters, and that we can rarely be *absolutely* certain of any of these things. But that just is to satisfy the conditions for the conflict diagnosis to kick in. Thus, this diagnosis seems to say that we're all guaranteed to be irrational *merely in virtue of our predicament as believers*.⁵⁵ That is a lot to swallow.

For my own part, I feel at least tempted to diagnose preface cases not as conflict cases but as what I earlier called *permissive* (counterexample) cases, where the thought is that substantive rationality *permits*, but doesn't require, us to believe the (highly probable) negation of the conjunction of one's beliefs – whereas structural rationality forbids having such a belief (while continuing to hold on to one's other beliefs). This still blocks the guarantee hypothesis, thus putting pressure on monism, but it does so without counting us as irrational merely in virtue of our predicament as believers. But it is kosher to adopt this diagnosis merely on the basis of wanting to avoid that latter result? That is a question for some soul-searching.⁵⁶

References

Baron, M. (2005). "Justifications and Excuses," *Ohio State Journal of Criminal Law*, 2/2: 387-406.

⁵⁵ This, it seems to me, is much bolder than saying that there are certain *very specific* cases where our circumstances conspire to guarantee our irrationality – as would be so if this occurred only when one is in a case of all-things-considered misleading higher-order evidence.

⁵⁶ For helpful discussion and comments, I'm grateful to members of my Fall 2019 seminar on structural rationality, especially Aaron Thieme, to Nick Hughes, and to two anonymous referees. This chapter was drafted before my 2021 book *Fitting Things Together: Coherence and the Demands of Structural Rationality* appeared, and many of the ideas and some exact passages were incorporated into the book, albeit in a different sequence. Thanks to Oxford University Press for permission to use the material in both places.

- Bergmann, M. (2005). "Defeaters and Higher-Level Requirements," *Philosophical Quarterly*, 55/220: 419-436.
- Broome, J. (2004). "Reasons," in Wallace, Pettit & Scheffler (eds.), *Reason and Value: Themes from the Moral Philosophy of Joseph Raz*. Oxford: Oxford University Press.
- Brown, J. (2018). *Fallibilism: Evidence and Knowledge*. Oxford: Oxford University Press.
- Chen, Y. & Worsnip, A. (forthcoming). "Disagreement and Higher-Order Evidence," in Baghramian, Carter & Cosker-Rowland (eds.), *The Routledge Handbook of the Philosophy of Disagreement*. Routledge.
- Christensen, D. (2004). *Putting Logic in its Place: Formal Constraints on Rational Belief*. Oxford: Oxford University Press.
- Coates, A. (2012). "Rational Epistemic Akrasia," *American Philosophical Quarterly*, 49/2: 113-124.
- Easwaran, K. & Fitelson, B. (2015). "Accuracy, Coherence, and Evidence," *Oxford Studies in Epistemology*, 5: 61-96.
- Enoch, D., Spectre, L. & Fisher, T. (2012). "Statistical Evidence, Sensitivity, and the Legal Value of Knowledge," 40/3: 197-224.
- Fantl, J. & McGrath, M. (2002). "Evidence, Pragmatics, and Justification," *Philosophical Review*, 111/1: 67-94.
- Fletcher, G.P. (1985). "The Right and the Reasonable," *Harvard Law Review*, 98/5: 949-982.
- Friedman, J. (2013). "Suspended Judgment," *Philosophical Studies*, 162/2: 165-181.
- González de Prado, J. (2020). "Dispossessing Defeat," *Philosophy & Phenomenological Research*, 101/2: 323-340.
- Harman, E. (2011). "Does Moral Ignorance Exculpate?," *Ratio*, 24/4: 443-468.
- Hawthorne, J., Rothschild, D. & Spectre, L. (2016). "Belief is Weak," *Philosophical Studies* 173/5: 1393-1404.
- Hawthorne, J. & Srinivasan, A. (2013). "Disagreement Without Transparency: Some Bleak Thoughts," in Christensen & Lackey (eds.), *The Epistemology of Disagreement: New Essays*. Oxford: Oxford University Press.
- Horowitz, S. (2014). "Epistemic Akrasia," *Noûs*, 48/4: 718-744.
- Kaplan, M. (1996). *Decision Theory as Philosophy*. Cambridge, UK: Cambridge University Press.
- Kelly, T. (2002). "The Rationality of Belief and Some Other Propositional Attitudes," *Philosophical Studies*, 110/2: 163-96.
- Kiesewetter, B. (2017). *The Normativity of Rationality*. Oxford: Oxford University Press.
- Kolodny, N. (2007). "How Does Coherence Matter?," *Proceedings of the Aristotelian Society*, 107/3: 229-263.
- (2008a). "Why Be Disposed to Be Coherent?," *Ethics*, 118/3: 437-463.
- (2008b). "The Myth of Practical Consistency," *European Journal of Philosophy*, 16/3: 366-402.
- Lasonen-Aarnio, M. (2020). "Enkrasia or Evidentialism? Learning to Love Mismatch," *Philosophical Studies*, 177/3: 597-632.
- (2021). "Coherence as Competence," *Episteme*, 18/3: 353-376.
- Lord, E. (2018). *The Importance of Being Rational*. Oxford: Oxford University Press.
- Lord, E. & Sylvan, K. (2021). "Reasons to Suspend, Higher-Order Evidence, and Defeat," in Simion & Brown (eds.), *Reasons, Justification, and Defeat*. Oxford: Oxford University Press.

- MacFarlane, J. (ms.). "In What Sense (if any) is Logic Normative for Thought?"
- Makinson, D.C. (1965). "The Paradox of the Preface," *Analysis*, 25/6: 205-207.
- Nelson, M. (2010). "We Have No Positive Epistemic Duties," *Mind*, 119/473: 83-102.
- Neta, R. (2018). "Evidence, Coherence and Epistemic Akrasia," *Episteme*, 15/3: 313-328.
- Ryan, S. (1991). "The Preface Paradox," *Philosophical Studies*, 64/3: 293-307.
- Schroeder, M. (2012). "The Ubiquity of State-Given Reasons," *Ethics*, 122/3: 457-488.
- Singh, K. (ms.). "Rationality and Correctly Responding to Reasons."
- Silva Jr., P. (2017). "How Doxastic Justification Helps Us Solve the Puzzle of Misleading Higher-Order Evidence," *Pacific Philosophical Quarterly*, 98/S1: 308-328.
- Skipper, M. (2019). "Higher-Order Defeat and the Impossibility of Self-Misleading Evidence," in Skipper & Steglich-Petersen, *Higher-Order Evidence: New Essays*. Oxford: Oxford University Press.
- Smithies, D. (this volume). "The Unity of Evidence and Coherence."
- Srinivasan, A. (2015). "Normativity Without Cartesian Privilege," *Philosophical Issues*, 25: 273-299.
- Thomson, J.J. (1986). "Liability and Individualized Evidence," *Law and Contemporary Problems*, 49: 199-219.
- Tucker, C. (forthcoming). "Withhold by Default: A Difference Between Epistemic and Practical Rationality," *Philosophical Studies*.
- Wedgwood, R. (2017). *The Value of Rationality*. Oxford: Oxford University Press.
- Williamson, T. (2000). *Knowledge and its Limits*. Oxford: Oxford University Press.
- (forthcoming). "Justifications, Excuses, and Skeptical Scenarios," in Dutant & Dorsch (eds.), *The New Evil Demon*. Oxford: Oxford University Press.
- Worsnip, A. (2016). "Belief, Credence, and the Preface Paradox," *Australasian Journal of Philosophy*, 94/3: 549-562.
- (2018). "The Conflict of Evidence and Coherence," *Philosophy & Phenomenological Research*, 96/1: 3-44.
- (2019). "Can Your Total Evidence Mislead About Itself?," in Skipper Rasmussen & Steglich-Petersen (eds.), *Higher-Order Evidence: New Essays*. Oxford: Oxford University Press.
- (2021). *Fitting Things Together: Coherence and the Demands of Structural Rationality*. New York: Oxford University Press.