

WHICH REASONS? WHICH RATIONALITY?

DANIEL FOGAL
New York University

ALEX WORSNIP
University of North Carolina at Chapel Hill

The slogan that rationality is about responding to reasons has a turbulent history: once taken for granted; then widely rejected; now enjoying a resurgence. The slogan is made harder to assess by an ever-increasing plethora of distinctions pertaining to reasons and rationality. Here we are occupied with two such distinctions: that between subjective and objective reasons, and that between structural rationality (a.k.a. coherence) and substantive rationality (a.k.a. reasonableness). Our paper has two main aims. The first is to defend dualism about rationality—the view that affirms a deep distinction between structural and substantive rationality—against its monistic competitors. The second aim is to answer the question: with the two distinctions drawn, what becomes of the slogan that rationality is about responding to reasons? We'll argue that structural rationality cannot be identified with responsiveness to any kind of reasons. As for substantive rationality, we join others in thinking that the most promising reasons-responsiveness account of substantive rationality will involve an “evidence-relative” understanding of reasons. But we also pose a challenge for making this idea precise—a challenge that ultimately calls into question the fundamentality of the notion of a reason even with respect to the analysis of substantive rationality.

1. Introduction

The idea that rationality is about responding to reasons has a turbulent history. For a time in the late 20th century, it seemed to be so dominant that it was often

Contact: Daniel Fogal <daniel.fogal@nyu.edu>
Alex Worsnip <aworsnip@unc.edu>

taken for granted without argument.¹ It then came under sustained assault,² and—at least in the practical rationality literature³—went into retreat. Yet more recently, the view has enjoyed a resurgence, with sophisticated versions of reasons-responsiveness theories of rationality abounding.⁴

One thing that makes it hard to assess the slogan that rationality is about responding to reasons is that there is an ever-increasing plethora of distinctions that are drawn not only with respect to reasons, but also with respect to rationality—yielding a corresponding plethora of interpretations of the slogan. Two such (purported) distinctions are especially important to us here: one with respect to reasons, and one with respect to rationality.⁵

The first is the distinction between **subjective** and **objective** reasons.⁶ Very roughly, subjective reasons are constrained by the epistemic situation or perspective of the agent for whom they are reasons, whereas objective reasons are not.

The second is the distinction between **structural** and **substantive** rationality.⁷ Very roughly, structural rationality is about whether one's attitudes *fit together* or *cohere* with each other, whereas substantive rationality is about whether one's attitudes are actually *reasonable* or *justified*.

While the first of these distinctions—or at least something in the neighborhood of it—is widely accepted, the second is controversial. Call the view that affirms that there are two distinct, genuine kinds of rationality—structural and substantive—*dualism* about rationality, and that which denies this *monism* about rationality.

The first of this paper's two main aims is to strengthen the case for dualism over its monistic competitors. It turns out that the distinction between different kinds of reasons—"objective" and "subjective"—bears on this aim. Some have thought that once we clarify what kind of reasons rationality requires us to respond to, we no longer need to distinguish substantive and structural rationality. We argue that this is not so. On the contrary: getting clear on the distinction

1. See, among many others, Foot (1972), Williams (1981).

2. See, among many others, Scanlon (1998: Ch. 1), Broome (2007; 2013), Setiya (2004).

3. Though not in epistemology, where it is typically taken as axiomatic that epistemic rationality consists in responding to evidence—which plausibly amounts to the view that epistemic rationality consists in responding to a particular kind of reason.

4. See, among others, Schroeder (2009), Gibbons (2010), Kiesewetter (2017), Lord (2018).

5. Various other distinctions won't occupy us. For example, we won't say anything about the distinction between *ex ante* and *ex post* rationality. Additionally, like other participants in the contemporary debate, we intend our discussion to cover both the rationality of doxastic states such as belief and the rationality of practical states such as intention.

6. See, among others, Schroeder (2004; 2007; 2009), Setiya (2004), Way (2009; 2012), Markovits (2014), Whiting (2014), Sylvan (2015).

7. See, among others, Scanlon (2007), Chang (2013), Neta (2015), Worsnip (2018), Fogal (2020).

lying behind the objective/subjective contrast actually *strengthens* the dualist position, and weakens the monist one.⁸

The second aim is to answer the following question: given dualism, and given the distinction between different kinds of reasons, what becomes of the slogan that rationality is about responding to reasons?

Our view is that *structural* rationality cannot be identified with responsiveness to *any* kind of reasons. This reinforces the depth of the dualism that we defend: it isn't that substantive rationality involves responsiveness to *one* kind of reasons, and structural rationality involves responsiveness to *another*—instead, structural rationality does not consist in responsiveness to reasons of any kind.

What about *substantive* rationality? It seems hard to deny that substantive rationality has at least *something* to do with reasons: after all, in glossing the difference between structural and substantive rationality, we said that substantive rationality is about being *reasonable*. Like some others, we think that the most promising interpretation of the slogan that substantive rationality is about responding to reasons will involve an “evidence-relative” understanding of reasons. But we also pose a challenge for making this idea precise—a challenge that ultimately, surprisingly, calls into question the fundamentality of the notion of a reason even with respect to the analysis of substantive rationality.

Since the dialectic and structure of this paper is complex, here is a guide. You can skip to §2 if you prefer the twists and turns of the journey to unfold dramatically.

In §2, we introduce in more detail the two distinctions we began with, via intuitive examples. §2.1 argues that the binary objective/subjective distinction is, as it stands, too simplistic, and that at least three notions need to be distinguished: fact-relative reasons, evidence-relative reasons, and belief-relative reasons. §2.2 turns to the substantive/structural distinction. Since this is the more controversial of the two distinctions, and its reality is at issue in this paper, it needs motivating. As such, we provide an initial case for it. This provides a background against which to evaluate the monist attempts to resist dualism that we consider later.

8. There are various objections to dualism that we won't discuss in this paper. For example, some (e.g., Kiesewetter 2017) argue that so-called structural rationality could not be normative, and that consequently it can't be a genuine form of rationality. We cannot address these other challenges for dualism in the space available here. (One of us does so elsewhere; see Worsnip in press: ch. 8.) We focus on making a positive case for dualism, on drawing attention to the inadequacy of prominent monist views, and on defending dualism against the specific charge that once we understand what kind of reasons rationality requires us to respond to, the structural/substantive distinction disappears.

§3 stays within the dualist framework, and asks the following question: if substantive rationality is about responding to reasons, which kind of reasons (fact-relative, evidence-relative, or belief-relative) does it require us to respond to? We argue that the fact-relative and belief-relative answers clearly fail, and that the evidence-relative answer is the most plausible.

§4 considers a range of monist rivals to dualism in light of our initial case for it. We can classify these different rivals according to whether they say that rationality *simpliciter* consists solely in coherence (Coherentist Monism), that it consists solely in reasons-responsiveness (Reasons-Responsiveness Monism), or that coherence and reasons-responsiveness come to the very same thing (Unificatory Monism). We also subdivide Reasons-Responsiveness Monist views into different categories depending on which kind of reasons they invoke: evidence-relative, belief-relative, or a hybrid of the two.⁹ We argue that all of these monist views fail.¹⁰

§5 considers a form of dualism that tries to understand substantive and structural rationality in terms of responsiveness to *different kinds* of reasons (evidence-relative and belief-relative, respectively). We argue that this view fails too. Taken together with our earlier arguments, this entails that structural rationality does not consist in responsiveness to any of the three categories of reasons we consider.

Finally, §6 returns to the one remaining interpretation of the claim that rationality consists in responding to reasons that still seems plausible: namely, that substantive rationality consists in responding to evidence-relative reasons. We raise a challenge for this view, and suggest the beginnings of a solution to it. However, that solution in turn leads us to doubt whether the notion of a reason really is fundamental to the analysis of substantive rationality.

2. The Two Distinctions

2.1. *The Objective/Subjective Distinction*

The classic case used to motivate the objective/subjective reasons distinction is Bernard Williams's (1981) famous petrol case. In this case, you have a glass

9. In principle, there could be a Fact-Relative Reasons-Responsiveness Monist view, but as we note briefly when we come to it, this view is deservedly unpopular.

10. This does not rule out every logically possible monist view, since there could be some monist views that draw on neither the notion of coherence nor that of reasons-responsiveness, or that try to unify rationality around a notion of reasons that is neither fact-relative, evidence-relative, belief-relative, nor some hybrid thereof. It does, however, rule out the main monist views in the literature.

of liquid that you believe to be gin and tonic, but in fact it is petrol. In the objective sense, it is said, you have a strong reason *not* to drink—since the liquid is petrol. But in the subjective sense, it is said, you have a strong reason *to* drink—since you believe the liquid to be gin and tonic (and you want a gin and tonic).

In fact, however, the dichotomy between “objective” and “subjective” reasons is insufficiently fine-grained.¹¹ Distinguish two versions of the petrol case:

Cleverly Disguised Petrol. Though it is in fact petrol, the stuff in the glass looks like gin and tonic, smells like gin and tonic, and has been served to you by a barman in response to your request for gin and tonic. On the basis of this evidence, you believe that it is gin and tonic.

Obviously Petrol. As well as in fact being petrol, the stuff in the glass looks like petrol, smells like petrol, and is sitting around in a car mechanic’s garage. In defiance of all this evidence, however, you believe that it is gin and tonic.

In *Cleverly Disguised Petrol*, your belief that the stuff in the glass is gin and tonic is justified by your evidence. In *Obviously Petrol*, it isn’t. This helps us to see that there are (at least) *three* bodies of information to which our reasons-talk can be relativized.

On a maximally objective, *fact-relative* notion of a reason, you have a strong reason to refrain from drinking both in *Cleverly Disguised Petrol* and in *Obviously Petrol*, since the glass in fact contains petrol in both cases; you do not have any comparably strong countervailing reason to drink.

On a maximally subjective, *belief-relative* notion of a reason, you have a strong reason to drink both in *Cleverly Disguised Petrol* and in *Obviously Petrol*, since in both cases you *believe* that the glass contains gin and tonic.

However, there seems to be an important notion of a reason intermediate between these two notions. On this intermediate notion of a reason, you have a strong reason to drink in *Cleverly Disguised Petrol*, but you lack a strong reason to drink in *Obviously Petrol* (indeed, you have a strong reason *not* to drink in *Obviously Petrol*). This notion of a reason is sensitive to the salient difference between *Cleverly Disguised Petrol* and *Obviously Petrol*, namely that in *Cleverly Disguised Petrol*, your belief that the glass contains gin and tonic is justified by

11. Feldman (1988: 407–8) astutely makes this point too.

your evidence, whereas in *Obviously Petrol*, it isn't. So let's call it an *evidence-relative* notion of a reason.^{12, 13}

We can summarize the differences between the three notions of a reason as follows:

	What do the reasons support in. . .	
	. . . <i>Cleverly Disguised Petrol</i> ?	. . . <i>Obviously Petrol</i> ?
Fact-relative (“objective”)	Refraining	Refraining
Evidence-relative	Drinking	Refraining
Belief-relative	Drinking	Drinking

The literature has been muddled by the fact that some philosophers use ‘subjective reason’ to mean (something like) ‘evidence-relative reason’,¹⁴ while others use it to mean ‘belief-relative reason’,¹⁵ and others either aren't clear which of

12. In distinguishing belief-relative and evidence-relative reasons, we are assuming that one's evidence does not consist in all of one's beliefs (including those that one holds unjustifiably). This view is rightly unpopular in recent epistemology. To borrow an example from §5 below, suppose that George (unjustifiably) believes that Obama is a Muslim and all Muslims are terrorists. If all George's beliefs were part of his evidence, he would have very strong (since entailing) evidence that Obama is a terrorist. But this result is obscene.

There are many other accounts of evidence available, but we won't be taking a stand on the correct account here. Instead, we'll assume that one can get a rough, adequate-enough grip on the distinction between evidence-relative and belief-relative reasons through cases like *Obviously Petrol*.

13. Note that an *evidence-relative* reason is not necessarily an *evidential* reason. An evidential reason for believing *p* is something that is evidence for *p*. It is unclear how to extend this to make sense of the notion of an evidential reason for action, intention, desire, preference, and so on. However, we can sensibly talk about what you have most practical (viz. moral, prudential, aim-given, or whatever) reason to do *given your evidence*. You have an evidence-relative *prudential* (or, perhaps, aim-given) reason not to drink in *Obviously Petrol* (roughly) because your evidence indicates that the glass contains petrol, and so that drinking will be bad for you. Similarly, you have an evidence-relative prudential reason to drink in *Cleverly Disguised Petrol* (roughly) because your evidence indicates that the glass contains gin and tonic, and so that drinking will bring you pleasure, or satisfy your desires. It may be that when it comes to reasons for *belief*, evidence-relative reasons are identical to evidential reasons. But even that is a substantive claim.

14. E.g., Gibbons (2013: Ch. 2), Markovits (2014: 7), Whiting (2014: 6–7), Wodak (2019: 237).

15. E.g., Schroeder (2004; 2007; 2009), Setiya (2004: 271), Portmore (2011: 14–15), Vogelstein (2012).

the two readings they intend¹⁶ or deliberately amalgamate the two.^{17, 18} We don't think these philosophers are best understood as giving competing analyses of a single, unified, pretheoretical target, the notion of a subjective reason. Instead, we take them to be picking out distinct notions that are worth distinguishing from each other just as much as they are worth distinguishing from fact-relative (or "objective") reasons. Consequently, we propose jettisoning the objective/subjective terminology in favor of a tripartite distinction between fact-relative, evidence-relative, and belief-relative reasons.^{19, 20}

2.2. *The Substantive/Structural Distinction* (*a.k.a., an Initial Case for Dualism*)

Let's turn now to the substantive/structural distinction. As we have already said, the reality of this distinction is controversial. But we can make an initial case for the reality of the distinction—that is, for dualism—by focusing on cases in which it seems that agents make two separate rational mistakes, of intuitively different kinds.²¹

Consider someone—let's call him Tom—who believes that he is Superman, and believes that Superman can fly. In believing that he's Superman, Tom has a belief that flies in the face of his evidence (or so we may safely stipulate). Moreover, it's very natural to describe this belief as irrational. But now let us add

16. E.g., Way (2012).

17. E.g., Schroeder (2011), Sylvan (2015). We'll consider this view below in §4.3.3.

18. To further add to the confusion, some epistemologists use 'objective reasons' to refer not to fact-relative reasons but to evidence-relative reasons (while using 'subjective reasons' to refer to belief-relative reasons). Feldman (1988: 411) interprets Pollock, Goldman, and Alston as all talking this way. Consequently, the term 'objective reason' is used by some philosophers to refer to the very same thing that some other philosophers use 'subjective reasons' to refer to—*viz.*, evidence-relative reasons.

19. Another advantage of this switch in terminology is that it avoids the confusion engendered by the fact that the language of 'subjective' and 'objective' reasons is sometimes (e.g., by Parfit 2011) used to mark an entirely different distinction, namely that between desire-given and desire-independent reasons.

20. In drawing the tripartite distinction, we don't mean to endorse the claim that the count noun 'reason(s)' in its broadly normative sense is semantically ambiguous. The distinction is compatible with—and, indeed, we are sympathetic to—a semantic theory on which the term is context-sensitive, with a contextually-determined parameter for background information. (Cf., among others, Henning 2014; Finlay 2014: Ch. 4; Wedgwood 2017: Chs. 4–5.) On this view, 'fact-relative', 'evidence-relative' and 'belief-relative' just draw attention to the nature of the information relative to which a given 'reason(s)'-claim is to be evaluated. Additionally, we leave open the possibility that there are further notions of a reason beyond the three we distinguish. These three, however, are the ones most relevant for our purposes.

21. The examples below are adapted from Fogal (2018). The Superman example originated with Jim Pryor.

another piece of information about Tom: he believes he can't fly. (Suppose he's tried several times, without success.) Intuitively, given this further piece of information, we can now see that Tom is irrational in a second respect: specifically, he fails to believe an obvious consequence of his other beliefs.

Note, to start with, that the second rational failing really does seem like a *second* rational failing, distinct from the first. We already know that Tom is irrational when we know that he believes that he is Superman, and believes that Superman can fly. But when we add that he believes he can't fly, Tom seems to be subject to an additional, new charge of irrationality. Tom is irrational in two respects, not one. Call this a *counting intuition* about rationality.

Moreover, the second rational failing seems importantly different in kind from the first. The first rational failing consists in Tom's failing to respond correctly to his evidence (by believing that he is Superman). But when we learn that Tom also believes that he can't fly, we do not learn that he has, in some second way, failed to respond correctly to his evidence. On the contrary, in fact: Tom's evidence *supports* believing that he can't fly.

Thus, it seems that we have two distinct kinds of irrationality. We can call these substantive and structural irrationality, respectively. Tom's substantive irrationality consists in his failure to respond correctly to his evidence. Since evidence for p is a paradigm instance of a reason to believe p , the natural generalization of this is to say that substantive rationality (as it applies to all attitudes—including not just beliefs but also intentions, preferences, and the like—as well as, perhaps, to action) consists in responding correctly to reasons, or being reasonable. Tom's structural irrationality consists in his having inconsistent beliefs. Since having inconsistent beliefs is a paradigm instance of incoherence, the natural generalization of this is to say that structural rationality (as it applies to all attitudes) consists in being coherent.

To further strengthen the case for distinguishing two different notions of rationality here, we can compare Tom to his brother Tim. Like Tom, Tim believes that he is Superman and that Superman can fly, but unlike Tom, Tim believes that he *can* fly. There's a clear sense in which Tim is even less rational than Tom, since he has two beliefs that go dramatically against his evidence (viz. that he is Superman, and that he can fly), where Tom has only one (viz. that he is Superman). But there is also a clear sense in which Tim is *more* rational than Tom, since his beliefs cohere in a way that Tom's don't: there's no inconsistency in his beliefs (and, indeed, his belief that he can fly is the logical consequence of his other two beliefs). But, it seems, we can only respect both of these points—marking both the sense in which Tim is less rational than Tom, and the sense in which he is more rational than Tom—by pulling substantive and structural rationality apart.

The distinction between structural and substantive rationality can be motivated using other forms of incoherence besides deductive inconsistency,

including incoherence between non-doxastic states. Consider an analogous case involving means-end incoherence (following Setiya 2007, inspired by Rawls 1971). Talia intends to count the number of blades of grass in her garden, even though doing so brings her no great pleasure, serves no other worthwhile purpose, and comes with a significant opportunity cost. In light of these facts, Talia lacks sufficient reason to count the number of blades of grass in her garden, and it's natural to describe her as irrational for intending to do so.²²

But now, again, let us add a further piece of information. Though Talia knows that in order to complete the count, she must keep track of how many blades she has counted so far, she can't be bothered to keep track and so doesn't intend to. Again, adding this information seems to involve Talia in a second kind of irrationality. But again, the problem isn't that she has good reason to keep track of the number of blades of grass she's counted, but fails to intend to do so. Since counting the blades of grass isn't a worthwhile project in the first place, Talia has no good reason to keep count of how many blades she has counted. Rather, the problem seems to be a distinctively different one of failing to intend the means to her ends. Again, we can call these distinct failures substantive and structural irrationality, respectively. Of course, there are other terms one could use to capture this distinction. But what matters is the reality of the distinction, not what we call it.

3. Substantive Rationality and Kinds of Reasons

In introducing the substantive/structural distinction, we glossed substantive rationality as a matter of being reasonable or, what seems to be equivalent, correctly responding to reasons. But with the distinction between fact-relative, evidence-relative, and belief-relative reasons on the table, we can ask: which notion of reasons is at work in this gloss?

The objective (or fact-relative) notion won't do, as we can see from *Cleverly Disguised Petrol*. In this case, your fact-relative reasons on balance support *not* drinking, since the liquid is petrol. There is hence a sense in which you ought not drink. But there is *no* ordinary sense of the term 'irrational' such that you are irrational in drinking in this case, when you couldn't possibly have known that the liquid was petrol. In any ordinary sense of 'irrational', you cannot be irrational for failing to respond to facts completely outside of your ken. Of course, a philosopher could stipulatively define a sense of 'rationality' on which

22. At least if she hasn't been exposed to highly misleading evidence that counting blades of grass is a worthwhile pursuit. See §3 below.

it requires responsiveness to all one's fact-relative reasons.²³ But this would not correspond to (and would risk obscuring) the central and ordinary concept of substantive rationality as we understand it: namely, a notion of rationality that requires more than just attitudinal coherence, but less than omniscience and less than responsiveness to considerations one could not possibly have known about.²⁴

On the other hand, we shouldn't understand substantive rationality in terms of responsiveness to the belief-relative reasons either. We can see this from *Obviously Petrol*. In this case, your belief-relative reasons on balance support drinking, even though all of your evidence indicates that the liquid in the glass is petrol. Now, you might (intentionally) drink in this case without any *structural* irrationality, or incoherence. But there is a clear and ordinary sense of 'irrational' in which it would be irrational to drink in this case. That is the sense that talk of substantive rationality (as distinguished from structural rationality, or coherence) should capture.

This leaves the evidence-relative notion of a reason as the best candidate for filling in the slogan that substantive rationality concerns responsiveness to reasons. On the evidence-relative notion of a reason, your reasons support drinking in *Cleverly Disguised Petrol*, but refraining from drinking in *Obviously Petrol*. This seems to yield the right results about substantive rationality in each case: it is substantively rational to drink in *Cleverly Disguised Petrol*, but not in *Obviously Petrol*.

Unsurprisingly, then, recent sophisticated accounts of substantive rationality have zeroed in on the evidence-relative notion, or something close enough.²⁵ Indeed, the evidence-relative view has long been the default view of (substantive) rationality in epistemology. If substantive rationality were to be understood in

23. Rawls (1971: 417), Gert (2004: 5–9, Ch. 7), and Portmore (2011: 14–15, 161–62) all define what they call "objective rationality" this way.

24. Indeed, this is exactly what happens to Portmore. He defines objective rationality in terms of fact-relative reasons, and he considers a definition of "subjective rationality" either in terms of coherence or in terms of belief-relative reasons, but he never considers the possibility of defining a notion of rationality intermediate between these. Given the risk of obscuring the notion of substantive rationality by using the term 'objective rationality', we think it is clearer to replace the latter with talk of what one objectively ought to do, or what one has most objective reason to do. This does less violence to ordinary usage, and doesn't seem to lose any expressive power.

25. Kiesewetter (2017) defends the evidence-relative view, whereas Lord (2018) defends the view that rationality requires you to respond to the reasons you possess, where one condition on possessing a reason is that you are in a position to know it. Lord's view can still be thought of as an evidence-relative view, since (i) what you're in a position to know plausibly depends primarily upon your evidence, and (more controversially, but still plausibly) (ii) some proposition *p* is part of your evidence only if you're in a position to know *p*. Both Kiesewetter and Lord, however, offer their views as accounts of rationality *simpliciter*, thinking they can eliminate structural rationality. We argue against their attempt to do this in §4.3.2.

terms of fact-relative reasons, we'd be rationally required to believe all and only the truths. If it were understood in terms of belief-relative reasons, we'd count as substantively rational just for having beliefs that we *believe* to be supported by our evidence, even if we are wrong or unjustified in thinking that. Neither of these views about (substantive) rationality is popular in epistemology. Rather, the most common view is that substantive rationality requires us to believe what our evidence supports. (This is the view, for example, that's at work in describing Tom's belief that he is Superman as irrational.) That is what the evidence-relative view of reasons predicts.

To be sure, there are challenges for making an account of substantive rationality in terms of evidence-relative reasons precise. We'll consider some of these challenges in §6. But at the very least, the view seems superior to the fact-relative and belief-relative views, given the verdicts about cases like *Cleverly Disguised Petrol* and *Obviously Petrol* that each view yields.

4. Monist Resistance to Dualism

On the emerging dualist view that we've been arguing for, (i) substantive and structural rationality are distinct; (ii) substantive rationality consists, roughly, in responsiveness to evidence-relative reasons; and (iii) structural rationality consists in coherence. Our initial case for this stance serves as a background against which to consider different forms of monism that deny the distinction between substantive and structural rationality, or that try to unify the two under a single notion of rationality.

4.1. Coherentist Monism

The first version of monism is represented by John Broome (2007). Broome begins by noting that rationality doesn't require responsiveness to (what we're calling) fact-relative reasons. He argues for this on similar grounds to ours, namely that in cases like *Cleverly Disguised Petrol*, you fail to respond to fact-relative reasons without thereby being irrational. The lesson, he thinks, is that if there is any true claim in the ballpark of the slogan that rationality requires correctly responding to reasons, it must be one according to which "a rational person's response to reasons [is] filtered through her beliefs in some way" (Broome 2007: 353).

Broome then argues that the best interpretation of this thought is that it's part of being rational that, for all Φ , you intend to Φ if you believe that your reasons require you to Φ (2007: 359–61); in other words, that your intentions line

up with your beliefs about your reasons.²⁶ But this “enkratic” requirement is one of the standard, core requirements of *structural* rationality—it requires a kind of coherence, forbidding you from simultaneously believing that your reasons require you to *F* and yet not intending to *F*. However, it isn’t the *only* requirement of structural rationality: there are others, like avoiding means-end incoherence, avoiding contradictory beliefs, and so on. Thus, on Broome’s view, a notion of rationality *simpliciter* as coherence subsumes (without being exhausted by) the only truth in the ballpark of the claim that rationality requires correctly responding to reasons. No further notion of rationality is needed. This gives us:

Coherentist Monism. Rationality *simpliciter* consists in having attitudes that satisfy various requirements of coherence (including, but not limited to, aligning your intentions with your beliefs about your reasons).

The problem with Coherentist Monism, however, is that it fails to capture many cases that *are* naturally described as irrational. *Obviously Petrol* is an example. As we said in §3, there is clearly a sense in which in *Obviously Petrol* (unlike the agent in *Cleverly Disguised Petrol*) you are irrational for (intentionally) drinking. Though both cases are ones in which the stuff in the glass is in fact petrol, the relevant difference is that in *Cleverly Disguised Petrol* you lack evidence for this, whereas in *Obviously Petrol* you have lots of evidence for it. Broome’s view is insensitive to this difference. In neither case does your (intentionally) drinking involve a failure to align your intentions with your beliefs about your reasons: indeed, in both cases you (intentionally) do exactly what you believe that your reasons support doing. Nor do you seem to exhibit any other kind of incoherence. Thus, Broome’s view doesn’t deliver a verdict of irrationality in either case.

The sense in which you are irrational to drink in *Obviously Petrol* is exactly what the notion of substantive rationality is designed to capture: a kind of irrationality that can be present even when the agent doesn’t exhibit any incoherence. Broome misses the legitimacy of this notion of rationality due to his tacit assumption that, given that rationality doesn’t require responsiveness to *objective* reasons, it must require only responsiveness to one’s *beliefs* about reasons. This overlooks the possibility of more intermediate possibilities, such as that rationality requires responsiveness to one’s *evidence-relative* reasons.²⁷

26. On one interpretation of belief-relative reasons, where belief-relative reasons are just things you believe to be reasons, this is close to being equivalent to saying that rationality requires you to correctly respond to your belief-relative reasons. But this is not the only (nor, perhaps, even the most plausible) way of understanding belief-relative reasons.

27. See also Kiesewetter (2017: 161), Lord (2018: 23–24). It’s odd that Broome overlooks this possibility given that he claims that the ‘ought’ that expresses the central normative concept is

In a new paper (Broome in press), Broome notes that others have wanted to distinguish structural and substantive rationality. However, he denies that the ordinary English word ‘rational’ ever refers to (what others call) substantive rationality, and as such holds that we should use a different term for it that avoids the language of rationality entirely, reserving ‘rationality’ solely for structural rationality. He writes:

So far as I can tell, this [substantive] use of ‘rational’ is not historically justified. ‘Rational’ has never had this normative meaning in common English. [. . .] [Take] a case where the agent ought to do or intend something, but believes she ought not to. Suppose there is nothing irrational about her false belief; it is supported by good—though misleading—evidence. Presented with a case like this, would ordinary English speakers use ‘rational’ in [the substantive] sense? In this sense it would be rational for the agent to do or intend what she rationally believes she ought not to do or intend. Would any ordinary English speaker say that? [. . .] I very much doubt that any ordinary English speaker would say it. [. . .] So far as I can tell, the [substantive] sense is an invention of philosophers. [. . .] The new sense of ‘rational’ simply leads to confusion. Most philosophers who write about rationality intend to write about it as it is commonly understood. That is my intention. Given all this, we should eschew the [substantive] sense of ‘rational’, and I do. (Broome in press: 6)

The passage is odd in numerous respects. First, Broome assumes that in a case where you (objectively) ought to Φ , but rationally believe that you ought not to Φ , the substantive sense of ‘rational’ has it that rationality requires you to Φ . But this mistakenly assumes that those who distinguish substantive and structural rationality understand substantive rationality in terms of responding to one’s *objective* (fact-relative) reasons. If substantive rationality instead consists in responding to one’s evidence-relative reasons, then it is plausible that, if you *rationally* believe that you ought not to Φ , substantive rationality does *not* require you to Φ . Again, Broome is overlooking the evidence-relative notion of a reason (and a correspondingly evidence-relative notion of substantive rationality). We agree that the ordinary word ‘rational’ doesn’t refer to responsiveness to *objective* reasons. But it doesn’t follow that it never refers to anything more than coherence. Instead, it often refers to responsiveness to evidence-relative reasons.

Second, Broome refutes his own view without noticing it. Consider the sentence, “Suppose there is nothing irrational about her false belief; it is supported

“prospective”, not objective, where the former is relative to evidential probabilities (Broome 2013: 41). One might have thought he would endorse something similar about reasons.

by good—though misleading—evidence.” Here Broome assumes that the rationality of a belief is a matter of its being supported by good evidence. But given that a belief’s being supported by evidence is just a matter of its being supported by *reasons* (of a particular kind), this just is to employ a substantive notion of rationality with respect to belief. Thus, in arguing that it is unnatural to use ‘rational’ in a way that picks out substantive rationality, Broome himself uses ‘rational’ in exactly such a way. Moreover, his doing so reflects a very natural and widespread tendency to use ‘rational’ in its substantive sense when evaluating beliefs. When we say that the climate change-denier, or the flat-earther, or the person who believes in fairies at the bottom of her garden, is irrational, we are not (necessarily) saying that her beliefs are internally incoherent—which they may not be—but instead that they aren’t supported by her evidence, such that she lacks adequate reasons for them. As with *Obviously Petrol*, these cases show the need for a substantive notion of rationality that goes beyond coherence.

Finally, the fact that Broome feels the need to stipulate that the agent’s belief that she ought not Φ is supported by her evidence suggests that he recognizes that, if this belief were *not* supported by her evidence, it would not be so unnatural to describe her as rationally required to Φ . But this fact isn’t captured by Broome’s own theory. To capture it, we need to make room for a distinctive notion of substantive rationality, as responsiveness to evidence-relative reasons.

Objection: can coherentist monism accommodate evidential requirements?

In the preceding, we’ve been assuming that failures to respond to one’s evidence are not *ipso facto* instances of incoherence. This was crucial both to our contention that the coherentist monist cannot capture the sense in which you’re irrational to drink in *Obviously Petrol*, and to our contention that Broome is betraying his own theory when he admits that failures to believe what one’s evidence supports are irrational.

But this might be challenged, on the following grounds. Plausibly, you don’t have something as evidence unless you bear some particular mental state toward it. Or, at least, whether you have something as evidence supervenes on your mental states. But then, it seems, the irrationality of believing against your evidence involves a bad relation between your mental states. For example, in *Obviously Petrol*, you believe that the stuff in the glass is gin and tonic (and intend to drink it) despite having certain mental *experiences* as of the stuff in the glass smelling like petrol. Perhaps, then, there is a sense in which your mental states “clash” in such a case. And what is incoherence if not such a clash between your mental states?

We'll grant for the sake of argument that whether you have something as evidence supervenes on your mental states. However, not just any kind of broadly "bad relation" between one's mental states counts as incoherence of the sort that coherentist monists like Broome think makes for irrationality. First, remember that Broome needs a notion of coherence narrow enough to exclude responsiveness to reasons.²⁸ But evidence for p just is a certain kind of reason to believe p .²⁹ Indeed, an evidential consideration is the paradigm instance (and, according to some, the *only* instance) of a reason for belief. It is thus contradictory to claim that rationality doesn't require responsiveness to reasons, but does require responsiveness to evidence. And any notion of coherence that includes responsiveness to evidence will be one according to which rationality does, contra Broome, require responsiveness to at least one (very important) kind of reason.³⁰

Moreover, the requirement to believe what one's evidence supports is importantly different from the other coherence requirements that Broome endorses. These other requirements are *wide-scope*, in the sense that when one has the combination of mental states that the requirement forbids, one can come to satisfy the requirement by revising *any* one of the offending attitudes. And they are *schematic*, in the sense that they all pick out general patterns of incoherent attitudinal mental states, where this involves abstracting away, at least to some extent, from the content of the particular attitudes involved. For example, it's supposed to be incoherent to believe p and not- p , whatever ' p ' is; it's supposed to be incoherent to intend to Φ , believe that Ψ -ing is necessary for Φ -ing, but not intend to Ψ , whatever ' Φ ' and ' Ψ ' are; etc.

The requirement to believe what one's evidence supports is not wide-scope. For example, when one has experiences as of the glass containing petrol (and no reason to distrust these experiences), one is (arguably) required simply to believe that the glass contains petrol; the requirement to believe what one's evidence supports in such a case can't be satisfied by revising one's experiences (if such a thing is even possible). And, while the evidential requirement(s) can be stated in a way that is schematic (e.g.: "if your evidence decisively supports p , you're required to believe p "), in order to say which combinations of *mental*

28. Except in the very weak sense that it requires you to align your intentions with your *beliefs* about reasons.

29. Kelly (2007) argues that to say that E is evidence for p (in the epistemologically central sense of 'evidence') is to make a normative claim to the effect that E is a certain kind of reason to believe or lend credence to p .

30. Broome might try saying that he only meant that rationality doesn't require responsiveness to one's *practical* reasons. The key difference here, he might think, is that one's evidence supervenes on one's mental states, whereas one's practical reasons don't. But this again overlooks the notion of an evidence-relative (practical) reason. To whatever extent that it's plausible that one's evidence supervenes on one's mental states, it's equally plausible that one's evidence-relative *practical* reasons supervene on one's mental states.

states this requirement forbids, we'd need to say which mental states make it the case that your evidence decisively support p . But *this* (often) can't be captured schematically: when some experience (for example) makes it the case that your evidence supports p , this is often ineliminably due to the particular content of the experience and the particular content of p (and, perhaps, contextual features of one's situation) in a way that cannot be abstracted into any general pattern.

Thus, Broome cannot view failures to respond to one's evidence as (ipso facto) incoherent—which, given Coherentist Monism, means he can't view them as irrational either.

4.2. Unificatory Monism

It might be thought that the arguments just given target Broome himself narrowly. Couldn't there be other Coherentist Monists, who don't share all of Broome's theoretical commitments, who could embrace the idea that failures to believe what one's evidence supports are a kind of incoherence?

Since, as we argued, a failure to believe what one's evidence supports just is a failure to correctly respond to a particular kind of reason, such a view would have to give up on the claim that rationality requires coherence *rather* than responsiveness to reasons. It is thus best thought of as a kind of *unificatory* view:

Unificatory Monism. Rationality *simpliciter* consists in coherence, where coherence is understood broadly to include failures to respond to one's evidence-relative reasons.

This view is advanced by Ralph Wedgwood (2017: 11–12). Wedgwood claims that once we see that rationality requires responsiveness only to evidence-relative reasons, and that failures to respond to one's evidence-relative reasons are a kind of incoherence, the distinction between substantive and structural rationality turns out to be a “distinction without a difference”.

We think that a use of ‘incoherent’ that includes all failures to respond to one's evidence-relative reasons stretches the ordinary meaning of the term. This can be brought out by cases where it's unobvious exactly what one's evidence supports. For example, suppose it's August 2020 and you possess a body of evidence about who will be the US president in 2021. Presumably at least some credences will be inadequately supported by that evidence. Suppose, just to fix ideas, that given all the evidence you have, a credence of 0.95 that Biden will be president is inadequately supported. Does it follow that you would be *incoherent* if you had credence 0.95 that Biden will be president? We find it a stretch to say so. If you arrive at credence 0.95 because you simply *misassess* the

probative force of your evidence that Biden will be President, judging it to be stronger than it really is, then we think it would be very unnatural to describe you as incoherent—though you are, by hypothesis, less than perfectly *substantively* rational.

Suppose, though, that there is some sense in which having 0.95 credence would make you “incoherent.” Still, throwing such cases together with paradigm instances of structural irrationality seems to collapse distinct phenomena, obscuring important differences between them. Consider the difference between (a) someone who simply misassesses the force of her evidence that Biden will be president, judging that it supports credence 0.95 and forming credence 0.95 as a result, and (b) someone who herself judges that the evidence doesn’t support credence 0.95 that Biden will be president, but goes on to form this credence anyway. There seems to be an important difference between these two characters, and it is useful to be able to express that difference by saying that the former is at least *structurally* rational, while the latter is not. To say that these two errors are of a piece with one another seems like willful insensitivity to a distinction that, whatever language we use to capture it, is clearly there.

Thus, Wedgwood hasn’t really said anything to show that there isn’t a meaningful distinction between substantive and structural rationality, where the latter corresponds to a notion of coherence narrower than his more expansive one.³¹ Here’s one plausible constraint on this narrower notion of coherence: one shouldn’t count as incoherent when the “clash” between one’s mental states is only due to a failure to correctly grasp a *substantive relation of support* between some fact (or proposition, or experience) and some response. This holds even when this substantive relation of support is fairly obvious to reasonable people. Thus, the “clash” between (say) the experience as of the glass containing petrol and believing that the glass contains gin and tonic does not amount to incoherence—at least not in any sense of ‘incoherence’ that avoids collapsing distinct normative phenomena.

31. Wedgwood (2017: 12) simply states that “for the purposes of understanding rationality-as-coherence [i.e., structural rationality], we need to be understanding the notion of ‘coherence’ in a sufficiently broad and general way,” viz. one that includes failures to respond correctly to one’s evidence. Since Wedgwood is explicitly considering the dualist’s proposal here, and trying to argue that the dualist’s distinction doesn’t make sense, this seems to imply that even the dualist must understand coherence in this very broad way. But we cannot see why Wedgwood thinks this. On the contrary, it seems clear that in order to give the dualist’s distinction between structural and substantive rationality a fair shake, and to see the notion as rationality-as-coherence (i.e., structural rationality) that is of distinctive interest, we need to understand ‘coherence’ in a *narrow* way that clearly delimits it from substantive rationality. Wedgwood doesn’t say anything to show that this narrower way of understanding coherence isn’t a perfectly sensible one.

4.3. Reasons-Responsiveness Monism

Let's now turn to those who employ monistic notions of rationality as responding to reasons. A monistic notion of rationality as responsiveness to (all of) one's objective (fact-relative) reasons clearly fails, for the same reasons that it failed as an account of substantive rationality within a dualist framework. For example, it yields the result that the agent in *Cleverly Disguised Petrol* is irrational, which is clearly wrong. However, many philosophers have advanced monistic notions of rationality as responsiveness to "subjective" reasons.³²

This view is sometimes explicitly supposed to account for why it's irrational to have incoherent attitudes, since (it's thought) incoherent attitudes just are failures to respond to one's subjective reasons.³³ Consider Tom, for example. Tom's belief that he is Superman and that Superman can fly, it might be said, give him a decisive subjective reason to believe he can fly. So, in not believing this (and instead believing he can't fly), he's failing to respond to his subjective reasons. But there may also be failures to respond to subjective reasons that don't involve incoherence. Thus, the view is the mirror image of Broome's, holding that a notion of rationality *simpliciter* as responsiveness to subjective reasons—responsiveness subsumes (without being exhausted by) the claim that rationality requires coherence. This gives us:

(Subjective) Reasons-Responsiveness Monism. Rationality *simpliciter* consists in responding correctly to subjective reasons (where this requires, but is not exhausted by, various kinds of coherence).

However, in light of the discussion in §2.1, we can ask: what is meant by 'subjective reason' here? Does it refer to belief-relative reasons, evidence-relative reasons, or something else?

4.3.1. Belief-Relative Reasons-Responsiveness Monism

Suppose first that 'subjective reasons' refers to belief-relative reasons.³⁴ (This was the notion at work in the above explanation of how Tom's belief that he can't fly constitutes a failure to respond to his "subjective reasons".) This view fails for simple reasons—ones that mirror the reasons why it wasn't plausible to understand substantive rationality in terms of belief-relative reasons, within

32. See, e.g., Schroeder (2009), Parfit (2011), Way (2012), Gibbons (2010; 2013), Whiting (2014), Sylvan (2015; in press).

33. Cf., e.g., Schroeder (2009).

34. As it does in Schroeder (2009), Parfit (2011), and Way (2012).

the dualist framework. The problem, recall, is that such a view can't capture the clear intuition that you are (in at least one good sense) irrational to drink in *Obviously Petrol*.³⁵ Ironically, this problem is shared between Belief-Relative Reasons-Responsiveness Monism and Coherentist Monism.

4.3.2. Evidence-Relative Reasons-Responsiveness Monism

Suppose next that 'subjective reasons' refers to evidence-relative reasons.³⁶ We conceded earlier that a notion of rationality as responsiveness to evidence-relative reasons is promising as an account of *substantive* rationality (within a dualist framework). But for it to be a promising notion of rationality *simpliciter*, it needs to capture intuitions about the irrationality of incoherent states. It might seem initially unclear how it would do this. However, Benjamin Kiesewetter (2017) and Errol Lord (2018) both advance (roughly) the following hypothesis:

The Guarantee Hypothesis. For any incoherent set of states $\{S_1 \dots S_n\}$ fitting a pattern typically taken to be structurally irrational, it is guaranteed that at least one of $\{S_1 \dots S_n\}$ is insufficiently supported by the agent's evidence-relative reasons.³⁷

The idea is that whenever one has incoherent states, at least *one* of the incoherent states must be insufficiently supported by one's evidence-relative reasons, and thus (by the lights of the evidence-relative version of Reasons-Responsiveness Monism) be irrational. For example, when one believes p and believes not- p , at least one of these beliefs must be insufficiently supported by one's evidence. If this is right, Kiesewetter and Lord think, then a monistic notion of rationality as responsiveness to evidence-relative reasons can capture the intuition that the incoherent attitudes typically *taken* to be structurally irrational always involve irrationality, without conceding dualism.

We think that the Guarantee Hypothesis is false. However, establishing that would require a whole article of its own.³⁸ So we'll settle for a more modest point: even if the Guarantee Hypothesis is true, an evidence-relative version

35. The view will have similar problems with the climate change denier, the flat-earther, etc.

36. As it does in Gibbons (2010; 2013), Whiting (2014), and Kiesewetter (2017).

37. Before Kiesewetter and Lord, Kolodny (2007) advanced a similar thesis. However, Kolodny (i) doesn't frame the thesis in terms of *evidence-relative* reasons specifically, and (ii) doesn't use it to try to bolster a monist account of rationality in terms of responsiveness to evidence-relative reasons.

38. One of us argues against the Guarantee Hypothesis elsewhere: see Worsnip (in press: Ch. 3).

of Reasons-Responsiveness Monism cannot capture the “counting intuitions” highlighted in making our initial case for dualism in §2.2.

Consider again the case of Tom, who believes he’s Superman, believes Superman can fly, but believes he himself can’t fly. The Guarantee Hypothesis predicts that at least one of Tom’s incoherent attitudes is insufficiently supported by his evidence-relative reasons. And so it is: his belief that he’s Superman is insufficiently supported by his evidence-relative reasons.

However, recall, when we discover that Tom believes he can’t fly, we learn that he is irrational in a *further* way that we didn’t already know about when we only knew that he believes he’s Superman (and that Superman can fly). But even given the Guarantee Hypothesis, the evidence-relative version of Reasons-Responsiveness Monism can’t account for this “counting intuition”. Once we know Tom believes he’s Superman, we already know that he has at least one attitude that is insufficiently supported by his evidence-relative reasons. Does the discovery that Tom believes he can’t fly teach us that he has a second attitude that is insufficiently supported by his evidence-relative reasons? No—the belief that he can’t fly is actually *supported* by his evidence-relative reasons. Thus, the evidence-relative version of Reasons-Responsiveness Monism can’t say that when we learn that Tom believes he can’t fly, we learn that he is irrational in a further way additional to the irrationality of his belief that he’s Superman. It detects irrationality in Tom only once, in virtue of his belief that he’s Superman.³⁹ What’s more, this view has no way of capturing the sense in which Tim, Tom’s coherent twin, is *more* rational than Tom. For Tim has the same evidentially-unsupported belief as Tom—that he is Superman—plus another—that he can fly. On the evidence-relative monist view, there can be no sense of ‘rational’ in which Tim is more rational than Tom.⁴⁰

39. Admittedly, this view can say that the fact that someone has incoherent attitudes is an *indicator* that one of her attitudes must be irrational. So in some contexts it might tell us that this person is irrational, when we did not know this before. But in Tom’s case, we already knew that he was irrational when we knew that he believes he’s Superman.

40. There’s a different way of understanding the purported upshot of the Guarantee Hypothesis, on which it’s not intended to show that we can eliminate the notion of structural (ir)rationality, but rather to provide a (reductive) account of it. On this view, to be substantively irrational is to fail to respond to one’s evidence-relative reasons, and to be structural irrational is to have a combination of states that *guarantees* one has failed to respond to one’s evidence-relative reasons—that is, to have a combination of states such that there’s no possible world in which one has evidence-relative reasons that support them. This view makes structural irrationality a stronger property than substantive irrationality, and is compatible with saying that the former corresponds to a special kind of criticism—perhaps it’s worse to have attitudes that are *guaranteed* to go against one’s evidence-relative reasons than ones that merely *do* go against one’s evidence-relative reasons. It can then explain how Tom comes in for two distinct kinds of rational criticism (one for having a belief that fails to respond to his evidence-relative reasons, and one for having a combination of beliefs that *guarantees* a failure to respond to his evidence-relative reasons), and how Tim

4.3.3. Hybrid Reasons-Responsiveness Monism

Are there any further alternatives for the Reasons-Responsiveness Monist? One initially intriguing possibility is to try to define a *hybrid* notion of a subjective reason that includes both (what we've been calling) evidence-relative and (what we've been calling) belief-relative reasons.⁴¹ Advocates of this hybrid view typically take subjective reasons to include not just the contents of beliefs, but also those of perceptual experiences and other evidentially-relevant non-doxastic states, such as intellectual seemings or the deliverances of memory.⁴² Schroeder (2011: 204), for instance, takes subjective reasons to be the contents of one's "presentational attitudes", understood as attitudes that "present their content to their subject as being true". This includes belief and perceptual experience, but not (for example) desire, wonder, supposition, or assumption. Hybrid Reasons-Responsiveness Monism, then, would say that rationality *simpliciter* requires responsiveness to subjective reasons, on this hybrid notion of a subjective reason.

However, the hybrid view can't provide a satisfactory account of cases like *Obviously Petrol*. In this case, there's a clear sense in which you'd be irrational if you intended to drink, given your evidence that the stuff in the glass is petrol. But there also seems to be a clear sense in which you'd be irrational if you *didn't* intend to drink, given your belief that it's gin and tonic, and your strong desire for a gin and tonic. Now, according to the hybrid view, you have at least one subjective reason to intend to drink (provided by your beliefs), and at least one subjective reason not to intend to drink (provided by your experiences). These reasons then combine in some way to yield a verdict about what you have most all-things-considered subjective reason to intend to do, and thus, on the Hybrid Reasons-Responsiveness Monist view, what it's rational to intend to do.

is in one respect better off (since he doesn't have states that *guarantee* a failure to respond to his evidence-relative reasons in the relevant sense).

However, this view affirms that there are two different kinds of irrationality, substantive and structural, with different extensions. Thus, it's a dualist rather than a monist view, as we've defined those terms. But it's a version of dualism that links structural and substantive rationality quite tightly—effectively reducing structural rationality to substantive rationality plus the notion of a guarantee. A more robust form of dualism, which is what we incline toward, denies the possibility of such reductions (either of the substantive to the structural, or vice versa). It's beyond the scope of this article to argue for this robust dualism over the reductive kind, but again, see Worsnip (in press: Ch. 3) for arguments against the latter. Of course, the particular reductive form of dualism just mooted requires the truth of the Guarantee Hypothesis, so if the Guarantee Hypothesis is false (as we think it is), then the reductive view fails too.

41. Accounts of subjective reasons that fit this bill include those given in Schroeder (2011) and Sylvan (2015).

42. We assume that perceptual experiences are part of one's evidence or otherwise evidentially-relevant.

But consider what the possible verdicts could be. First, the reason to drink could outweigh the reason not to drink. If that's so, there's no explanation of why you'd be irrational if you intended not to drink (since it's not irrational to fail to do or be as an outweighed reason recommends). Second, the reason not to drink could outweigh the reason to drink. If that's so, for parallel reasons to the first case, there's no explanation of why you'd be irrational if you intended to drink. Or third, it could be that neither reason outweighs the other. But when neither of two competing reasons outweighs each other, there's no requirement to do as *either* recommends. So then, there's no explanation *either* of why you'd be irrational to intend to drink *or* of why you'd be irrational to intend not to drink.⁴³

A different but related problem for Hybrid Reasons-Responsiveness Monism is that it seems odd to think of reasons provided by evidence and reasons provided by beliefs as weighing against each other in the first place. It doesn't seem that, in *Obviously Petrol*, the reason to drink and the reason not to drink are countervailing considerations that can be felt from the same perspective to jointly determine an all-things-considered verdict. Rather, it seems that from one point of view, you have a strong reason to drink, and that from a different point of view, you have a strong reason not to drink. This phenomenology differs from ordinary cases of competing *pro tanto* reasons.

Other cases bring out the problem even more sharply. Consider George, who believes (unjustifiably, it goes without saying) that Obama is a Muslim and that all Muslims are terrorists. Notwithstanding these beliefs, we can agree that George has strong reasons, provided by his evidence, against believing that Obama is a terrorist. But it just seems wrong to say that George's unjustified beliefs provide him with competing reasons of the same broad kind in favor of believing that Obama is a terrorist. It's not just that the evidential reasons in favor of the former outweigh or even defeat the latter.⁴⁴ Rather, it's that any sense in which his unjustified beliefs provide him with "reasons" to believe that Obama is a terrorist can only be a sense of 'reason' *so* subjective that it's detached from the substantive reasonableness or justification of his beliefs entirely. Along the dimension of evaluation that's concerned with the substantive justification or reasonableness of his belief that Obama is a terrorist, his beliefs that Obama is a Muslim and that all Muslims are terrorists don't help *at all*. To amalgamate his

43. Similar points apply in Tom/Tim-style cases. Consider someone who believes (without sufficient evidence) that he is Superman and that Superman can fly. There's a sense in which such a person is irrational (a second time over) if he believes he can't fly (as Tom does), given his other beliefs. And there's also a sense in which such a person is irrational (a second time over) if he believes he *can* fly (as Tim does), given his evidence. Again, the hybrid view can't deliver both of these verdicts.

44. Schroeder (2011) holds that unjustified beliefs provide reasons, but ones that are guaranteed to be defeated.

evidential reasons to believe that Obama is not a terrorist with his belief-given “reasons” to believe that Obama is a terrorist is to fail to keep distinct normative phenomena apart.⁴⁵

If the reasons provided by evidence and the reasons provided by beliefs don’t weigh against each other, and don’t contribute to the same overall normative status (viz. being substantively justified), then this strongly suggests that combining them into one hybrid notion of a subjective reason is a mistake. And if that’s so, then it can’t be used to provide a monist account of rationality.

5. Reasons-Responsiveness Dualism

In the last section, we argued against a range of monist proposals, on which there is just one notion, rationality *simpliciter*, rather than two distinct kinds of rationality, substantive and structural. However, there are some views that are dualist by this definition, but that still preserve the thought that rationality *as a whole* consists in responding to reasons. These are views that say that substantive rationality is about responding to one kind of reasons, while structural rationality is about responding to another kind of reasons. Call these views Reasons-Responsiveness Dualist views.

A naïve version of this view says that substantive rationality is about responding to “objective” reasons, while structural rationality is about responding to “subjective” reasons. But, at least if ‘objective’ means ‘fact-relative’, we’ve already seen (in §3) that this is wrong; for even substantive rationality does not require responsiveness to (all of) one’s fact-relative reasons. A more plausible version of the view is this:

Reasons-Responsiveness Dualism. Substantive rationality is about responding to *evidence-relative* reasons, while structural rationality is about responding to *belief-relative* reasons.

Indeed, this is a natural view for someone initially inclined toward Hybrid Reasons-Responsiveness Monism to adopt in light of our criticisms of it in §4.3.3. Our criticisms suggested that combining evidence-relative and belief-relative reasons into a single, hybrid notion of a subjective reason is a mistake.

45. As the scare-quotes indicate, perhaps belief-given “reasons” aren’t really reasons at all, in any sense. If so, we’d have to reject the belief-relative sense of ‘reason’. Dancy (2000: Ch. 3) argues that we should do this, and that requirements of structural rationality can do all the work that the notion of a belief-relative reason was supposed to do. We stay neutral on whether this is right. If belief-given reasons don’t exist, all the worse for attempts to understand rationality in terms of them, or in terms of a hybrid notion that includes them.

Reasons-Responsiveness Dualism acknowledges this, while preserving the hybrid view's thought that both evidence-relative and belief-relative reasons in some way bear on rational evaluations.⁴⁶

As we said in §3, it is plausible that substantive rationality consists in something like responsiveness to evidence-relative reasons (though see §6 for some significant complications). However, we deny that structural rationality consists in responsiveness to belief-relative reasons.⁴⁷

The cases of structural irrationality where Reasons-Responsiveness Dualism does best are those involving an attitude (or absence thereof) plus some belief that casts that attitude in a negative light: for example, intending to Φ while believing that you ought not Φ .⁴⁸ In this case, your intention to Φ seems to constitute a failure to respond correctly to the belief-relative reason provided by your belief that you ought not intend to Φ .

However, there are many other kinds of structural irrationality that do not seem to be constituted by failures to respond correctly to one's belief-relative reasons. In fact, many instances of structural irrationality don't involve any beliefs at all. For example, it's structurally irrational to have cyclical preferences—preferring A to B, preferring B to C, and preferring C to A. None of these states are beliefs, and so none of them constitute or supply any belief-relative reason not to have the others. Similarly for having inconsistent intentions—intending to Φ while also intending not to Φ —and for many other cases.

There are also examples of structural irrationality involving *doxastic* states that aren't beliefs. For example, it's structurally irrational to have credence 0.6 for p while having credence 0.5 for not- p . And it's plausibly structurally irrational to intend to Φ to suspending judgment on whether one will be able to Φ . Again, neither state here seems to be a failure to respond to a reason provided by a belief.⁴⁹

46. Notably, Kurt Sylvan, who advances a hybrid view of subjective (or “apparent”) reasons in Sylvan (2015), has more recently (Sylvan in press) proposed splitting apparent reasons into two categories, roughly corresponding to evidence-relative and belief-relative reasons. Even more notably, he calls the two categories ‘substantive apparent reasons’ and ‘structural apparent reasons’, strongly implicating that he thinks the former gives an account of substantive rationality and the latter an account of structural rationality. He thus may be an example of someone who has moved from Hybrid Reasons-Responsiveness Monism to Reasons-Responsiveness Dualism.

47. The reasons that we give for this in what follows all constitute further reasons to reject the Hybrid Reasons-Responsiveness view as well, since they also suggest that the hybrid view can't capture our intuitions about structural (ir)rationality.

48. Some other examples: having attitudes (such as hope, fear, or belief) that you believe you ought not have, having inconsistent beliefs, or failing to believe an obvious logical consequence of your other beliefs.

49. See also Wedgwood (2017: 99–103).

One might try to plug these problems by extending the notion of a belief-relative reason into that of an attitude-relative reason, where this includes not just beliefs but other doxastic states like credence and suspension of judgment, and perhaps even conative states like intention and preference. However, this is trickier to do than it might initially seem. Theorists of belief-relative (or “subjective”) reasons often want to say that one’s reason is constituted by the *content* of one’s belief, not by the belief itself.⁵⁰ This content *gets* to be a reason by being believed, but the content, not the belief, constitutes the reason. Saying this allows them to avoid saying that any reasons are mental states, and allows them to understand fact-relative, evidence-relative, and belief-relative reasons as all being constituted by the same sort of entity (such as a proposition). But it’s hard to see how this works for these other states. It doesn’t seem like having a credence of (say) 0.6 for p suffices for my having the content p as a reason (of any sort), nor that it suffices for having any other particular proposition as a reason.⁵¹ And it’s even less clear how this would go for suspension of judgment, and still less for intention or preference.

Another problem for Reasons-Responsiveness Dualism applies even in the cases of structural irrationality in which the view does *best*, such as that of intending to Φ while believing that you ought not Φ . Consider Huck Finn, who believes that he ought not help Jim the runaway slave to escape, but nevertheless intends to help him escape. Given Huck’s belief that he ought not help Jim escape, he has a strong belief-relative reason to abandon his intention to help Jim. All other things equal, then, Reasons-Responsiveness Dualism says that Huck is rationally required to abandon his intention to help Jim. But this seems questionable. According to broadly “wide-scope” views of structural rationality,⁵² structural rationality at least *permits* Huck to resolve his incoherent state by, instead, giving up his belief that he ought not help Jim escape. But this is not a way for Huck to do what his belief-relative reasons support doing—they support giving up his intention to help; hanging on to his intention and giving up his belief is not a way to do that. Thus, at least if one shares the wide-scooper’s sense that structural rationality permits Huck to give up his belief (rather than giving up his intention), Reasons-Responsiveness Dualism gets the wrong result here.⁵³

50. See, e.g., Schroeder (2009; 2011).

51. Schroeder (2018) argues that this can be resolved by adopting a sophisticated form of expressivism about epistemic expressions (such as claims about epistemic probability), on which it turns out, eventually, that propositions are not the objects of credence. We cannot engage this complex and highly revisionary proposal in the space available here. But even if it succeeds, it’s not clear how it generalizes to the other states, such as suspension of judgment and preference.

52. Famously associated with Broome (2013), among others.

53. One advantage of the wider-scooper’s view is that it avoids positing a conflict between what structural rationality requires of Huck and what substantive rationality and/or morality requires of him. Since substantive rationality and/or morality presumably require Huck to intend

6. Substantive Rationality as Responsiveness to Evidence-Relative Reasons

So far we've defended our dualist view against its monistic competitors and defended our further view that structural rationality does not consist in responsiveness to *any* kind of reason (not even belief-relative reasons). Where does that leave the slogan that rationality consists in responsiveness to reasons? Well, it entails that rationality does not *entirely* consist in responsiveness to reasons. Nevertheless, it seems hard to deny that *substantive* rationality is about responsiveness to reasons: after all, we glossed it in terms of reasons-responsiveness, or being reasonable. In §3, we argued that the most plausible notion of a reason to plug into this slogan is the evidence-relative one. Thus, it might seem like we have our answer: the one true disambiguation of the slogan that rationality consists in responsiveness to reasons is this: *substantive* rationality consists in responsiveness to *evidence-relative* reasons.

However, in this final section, we'll present a challenge for fleshing out the notion of an evidence-relative reason in such a way as to yield an adequate account of substantive rationality. Eventually, this will lead us to call into question the fundamentality of the notion of a reason even with respect to the analysis of substantive rationality. To set the stage, we'll consider and criticize a recent proposal due to Errol Lord (2018).

6.1. Lord's Practical Condition

The evidence-relative view of substantive rationality is naturally understood as imposing an epistemic constraint on the reasons relevant to rationality. Such a constraint is motivated by cases like *Cleverly Disguised Petrol*. There are different ways of explicating this constraint, but Lord's version of it is that for some reason R to be relevant to the rationality of A's Φ -ing (or, as Lord says, for A to *possess* this reason),⁵⁴ A must be in a position to know R.⁵⁵ Lord calls this the "epistemic" condition on possessing a reason.

to help Jim, any view of structural rationality that requires him to give up this intention is committed to there being a conflict.

54. Lord thinks that a reason bears on rationality iff it's possessed. Whether the conditions he gives for possessing a reason track the ordinary sense in which we talk of possessing or having a reason, though, is not important. What matters is whether these conditions pick out the reasons relevant to rationality.

55. See Footnote 25 on how this can still be understood as a broadly evidence-relative view.

According to Lord, however, this epistemic constraint isn't enough to capture all and only the reasons that are relevant to rationality. To show this, he appeals to the following case, adapted from Broome (2007):

Lois's Fish. Lois just ordered fish from her favorite seafood restaurant. Right before she digs in, the waiter comes out to inform her that the fish contains salmonella. Lois has the unfortunate belief that salmonella is one of the many bacteria found in food that is harmless to humans. And, indeed, this belief is rational. A renowned food scientist told her so. So she goes ahead and forms an intention to eat the fish and eats the fish. (Lord 2018: 98)

In this case, Lord assumes, the fact that the fish contains salmonella is a decisive reason not to eat it. Moreover, Lois is in a position to know that the fish contains salmonella.⁵⁶ So, if meeting the epistemic condition suffices for possessing a reason, Lois possesses a decisive reason not to eat the fish. But, Lord judges, it isn't irrational for Lois to eat the fish. Thus, so far, the account of possessed reasons doesn't form the basis of a satisfying account of substantive rationality.

Lord's solution is to introduce a second, "practical" condition on possessing a reason, and hence on a reason's being relevant to rationality. The practical condition says that for A to possess R as a reason to Φ , A must be "in a position to manifest knowledge about how to use R as a reason to Φ " (2018: 121). Lord clarifies how he understands the relevant know-how thus: "if you know how to use R as a reason to Φ , you are disposed to Φ when R is a reason to Φ " (2018: 121).

We will argue that the introduction of this second condition makes Lord's theory of substantive rationality far too underdemanding. To see why, note that Lord accepts the following three claims:

- (1) Substantive rationality requires A to Φ only if (i) there is a decisive reason (or body of reasons) R for A to Φ , and (ii) A possesses R . (*general account of substantive rationality*)
- (2) A possesses a reason R to Φ only if A is in a position to manifest knowledge about how to use R as a reason to Φ . (*practical condition on possessing a reason*)
- (3) A is in a position to manifest knowledge about how to use R as a reason to Φ only if A is disposed to Φ when R is a reason to Φ .⁵⁷ (*clarification of how to understand the practical condition*)

56. This distinguishes the case from one like *Cleverly Disguised Petrol*, where the agent is *not* in a position to know the relevant reason, namely that the glass contains petrol.

57. There is a scope ambiguity in this premise. ' A is disposed to Φ when R is a decisive reason (or body of reasons) to Φ ' might be read either narrow-scope—such that the disposition is to $[\Phi]$,

It follows from (1), (2), and (3) that

- (4) Substantive rationality requires A to Φ only if (i) there's some decisive reason (or body of reasons) R for A to Φ , and (ii) A is disposed to Φ when R is a reason to Φ .

But, whenever 4(i) satisfied, R is a reason to Φ . So whenever 4(i) and 4(ii) are *both* satisfied, A is disposed to Φ . Thus, (4) entails:

- (5) Substantive rationality requires A to Φ only if A is disposed to Φ .⁵⁸

But (5) is not tolerable; it makes the account of substantive rationality too underdemanding. Paradigm cases of substantive irrationality involve agents who are not even *disposed* to do what substantive rationality requires them to do.⁵⁹

Consider first paradigm examples of (substantive) epistemic irrationality, such as a flat-earthier. The flat-earthier, let's suppose, is aware of various considerations—such as the fact that all scientists hold that the earth is a sphere—that are, in fact (though not in the flat earthier's opinion), decisive evidential reasons to believe that the earth is a sphere. In defiance of this evidence, the flat-earthier believes that the earth is flat, rather than a sphere. On Lord's account, the flat-earthier is (substantively) irrational only if he is *disposed* to believe that the earth is a sphere (in the presence of the relevant evidence), but nevertheless ends up believing that the earth is flat (in the presence of the very same evidence). But this seems wrong. It's completely natural, in the (substantive) sense of 'irrational' dominant in mainstream epistemology,⁶⁰ to call the flat-earthier irrational even if he's a steadfast flat-earthier with no disposition whatsoever to believe that the earth is a sphere. Indeed, this insensitivity of the flat-earthier's dispositions to the evidence is plausibly part of what actually *grounds* the verdict of severe substantive irrationality.

and A has this disposition when R is a decisive reason to Φ —or wide-scope—such that the disposition is to [Φ when R is a decisive reason to Φ], and A has this disposition whether R is a decisive reason to Φ or not. Whichever way it is read, though, the problematic consequence remains; see the next footnote below.

58. If premise (3) above is read wide-scope (see the last footnote), then perhaps strictly speaking (5) should be replaced by: (substantive) rationality requires you to Φ (in circumstance C) only if you are disposed to [Φ in circumstance C]. But this result is just as objectionable.

59. Sylvan (2015) includes a similar practical condition on the reasons that are "apparent" to you, so his view faces an analogous problem.

60. Note that Lord (2018: 4–5) claims to be siding with the epistemologist's ordinary conception of rationality, as compared with the (coherentist, Broomean) conception that was (at least for a time) dominant in the practical rationality literature.

Now consider cases of substantive practical irrationality, such as the agent who drinks from the glass in *Obviously Petrol*. It isn't part of the setup of this case that the agent is *disposed* to refrain from drinking, but goes on to drink anyway. Rather, we can understand the agent as having no disposition to refrain from drinking whatsoever, despite all the signs that the glass contains petrol. Once again, this lack of disposition to refrain from drinking doesn't mitigate the agent's substantive irrationality for drinking; if anything, it intensifies it.

The problem can be brought out further by considering a case that forms a minimal pair with *Lois's Fish*:

Lewis's Fish. Lewis orders fish from his favorite seafood restaurant. Right before he digs in, the waiter comes out to inform him that the fish contains salmonella. Lewis has the unfortunate belief that all bacteria, including salmonella, are harmless for humans, because they are "natural". This belief is not supported by Lewis's evidence: he has had previous experiences getting sick from eating spoiled food, and is aware of the scientific consensus that some bacteria are harmful. Nevertheless, it's what he believes. So he goes ahead and forms an intention to eat the fish and eats the fish.

Though we agree with Lord that *Lois* is rational to eat the fish, it is incredibly natural to say that *Lewis* is (substantively) irrational to eat the fish. But Lord's theory seems to give the same verdict about both. In both cases, it's natural to suppose, the agent is not disposed to treat the fact that the fish contains salmonella as a reason to refrain from eating the fish. Thus, on Lord's view, s/he doesn't possess this reason. Thus, on Lord's view, neither *Lois* nor *Lewis* is irrational. Intuitively, the difference between the two cases is that the background belief that the action rests on (namely, that salmonella is harmless to humans) is rational as held by *Lois*, but irrational as held by *Lewis*. However, this doesn't make a difference to the verdict that Lord's theory yields about the rationality of the action.

It might be suggested that this is OK as long as we can say that *Lewis* is irrational in *some* way: if not in his action, then in his background belief.⁶¹ But Lord's view can't even say that *Lewis's* background belief is irrational. For just as in the case of the flat-earthier, Lord's view will say that *Lewis* is irrational to believe that salmonella is harmless to humans only if he is *disposed* to treat the experiences of sickness and scientists' testimony as reasons not to believe that salmonella is harmless for humans—which, we can stipulate, he isn't. Thus, although Lord stipulates that *Lois* has good reasons for her background belief in

61. Compare Lord (2014: 161–62).

the original case, it's not clear why this should make a difference to Lois's rationality on Lord's theory.

That said, a similar problem for Lord's view arises even in cases that (unlike *Obviously Petrol* and *Lewis's Fish*) don't involve any intuitively irrational background beliefs. For instance, return to the example of Talia, whose primary aim is to count blades of grass, even though this doesn't bring her any pleasure. It's natural to describe even an ideally coherent version of Talia as seriously irrational, simply in virtue of her commitment to such a worthless end. In this case, her intention to count blades of grass need not rest on any intuitively irrational belief. Once again, Lord's theory would only be able to count Talia as irrational if she has some *disposition* not to count grass that is sensitive to her reasons not to count grass (but continues to count grass anyway). But again, it seems that we're willing to convict Tania of substantive irrationality whether or not this is so.

Lord (personal communication) has suggested, in reply, that the cases have to look very odd for the subject to genuinely have no disposition of the relevant kind, and that his account can still say that the subjects are irrational in all ordinary, non-far-fetched versions of the cases. We have three counter-replies. To make them concrete, we'll focus on the flat-earthier, though the same points apply to the other cases.

First, it doesn't seem that far-fetched to imagine the flat-earthier having no disposition of the relevant kind. We can imagine the flat-earthier as a conspiracy theorist who takes the testimony of scientists to be evidence *against* the content of what they say. Since this person is disposed to become *less* confident that the earth is a sphere in response to scientific testimony to that effect, it's plausible that he has no disposition to believe that the earth is a sphere in response to this testimony. This doesn't seem far-fetched or even that unusual. People can be (grossly) mistaken about what their evidence supports. There isn't always a part of them that recognizes what it really supports, with some corresponding disposition to respond appropriately, buried deep down. Evidential support relations just aren't that transparent.

Second, even if the case is unusual, it is possible, and Lord's view still seems to yield the wrong verdict about it, in a way that reveals a more general structural flaw with the view. Intuitively, the less that one is disposed to respond appropriately to one's evidence, the *more* (substantively) irrational one is, with the limiting case being that where one has no disposition whatsoever to respond appropriately to one's evidence, which is the *height* of (substantive) irrationality. But Lord's view predicts that having no disposition whatsoever to respond appropriately to one's evidence should let one off the hook. That is the wrong result.

Third, suppose Lord can construe the case so that in any realistic version of it, the agent has *some* disposition to believe that the earth is a sphere in response

to scientific testimony to that effect. In order for this to help Lord, even this very weak disposition (that is defeated by a stronger disposition to respond in the opposite way to *the very same stimulus*) needs to suffice for meeting the practical condition. In other words, he must understand his practical condition as requiring only *some* disposition, no matter how weak or defeated. But then Lord owes us an explanation of why, if we should expect there to be a very weak disposition of this kind in the flat-earther case, we shouldn't also expect there to be a similarly weak disposition for Lois to appropriately respond her reason not to eat the fish in *Lois's Fish*. (Recall that Lord wanted to say that Lois does *not* have this disposition, so as to say that she doesn't meet the second condition for possessing this reason, so as to say that she's not irrational to eat the fish.) But we see no principled grounds for holding this. If we're understanding dispositions weakly enough to guarantee that the flat-earther has some disposition to respond correctly to her reason to believe the earth is a sphere, why not also say that Lois has some disposition to respond correctly to her reason not to eat the fish?

Thus, Lord's view faces a dilemma: either understand dispositions, and the strength of them that's required, such as to count *both* the flat-earther *and* Lois as satisfying the practical condition (in which case the practical condition doesn't do the work it was supposed to do, viz. get Lois off the hook for irrationality), *or* understand them in such a way as count neither of them as satisfying the practical condition (in which case it can't count the flat-earther as irrational).

The upshot is that Lord's account fails to capture the irrationality in the very sorts of cases that motivate the introduction of a category of *substantive* rationality—one that goes beyond mere coherence—in the first place. Lord doesn't himself employ distinct notions of substantive and structural rationality. But he does seem to think that the idea that rationality consists *only* in coherence will be too underdemanding a theory of rationality. However, as we've seen, given the strong constraints Lord puts on what it takes to "possess" a reason, his own theory also ends up being similarly underdemanding. On Lord's view, in effect, the only way you can be (substantively) irrational in not Φ -ing is if you are *disposed* to Φ , but this disposition is somehow blocked. It's an oddly restrictive view of substantive rationality that confines it to failures to manifest dispositions that you already have.

6.2. The General Challenge

We've just argued that Lord's version of the evidence-relative view provides an inadequate account of substantive rationality. Although Lord's view is just one possible version of this view, its failure sets a more general task for other accounts. Lord is right that we don't want our account of substantive rationality to convict

Lois of irrationality. But, we have urged, we *do* want our account of substantive rationality to convict *Lewis*—and other characters like the flat-earther, the agent in *Obviously Petrol*, and so on—of irrationality. The task for evidence-relative theories is to find a view that satisfies both of these desiderata. It bears stressing that *Lois* *does* appear to meet Lord's first, epistemic condition on what it takes to possess a reason: she is in a position to know the fact (or proposition) that is a reason to refrain from eating (namely, that the fish contains salmonella), and that fact is plausibly part of her evidence. So the solution can't just be to reject Lord's second condition and hang on to the first. There was something real that motivated the introduction of the second condition.

6.3. *The Beginning of a Solution?*

How is this problem to be handled? Here's (the beginning of) a proposal. In the discussion so far, we've acceded to the common—and usually tacit—assumption that normative support for an agent's responses is (typically) fundamentally provided by atomic facts, taken on their own. For example, we've followed Lord in assuming that in *Lois's Fish*, the normative support for a certain response (viz. not eating the fish) is provided by the atomic fact that the fish contains salmonella. On this picture, this atomic fact enjoys a privileged status as the reason (where 'reason' here is a count noun) to refrain from eating the fish.⁶² Other facts, like the fact that salmonella is harmful to humans, are not themselves part of the reason, but merely help *explain why* the fact that the fish contains salmonella is a reason for *Lois* to refrain from eating it. Call this the *atomic view*. So long as we assume the atomic view, it seems that *Lois* meets the relevant epistemic condition for possessing the reason: she's in a position to know (indeed, does know) that the fish contains salmonella.

But the atomic view isn't the only option. According to what we'll call the *cluster view*, normative support is (typically) fundamentally provided by *clusters* of facts, and the atomic facts we typically cite as normative reasons serve as *representatives* of such clusters.⁶³ The claim here is not just the banal one that there's often more than one reason to do something (as when the fact that the fish contains salmonella is one reason to refrain from eating, and the fact that the fish is over-salted is another). Rather, it's that even a *single* strain of normative support is typically provided by a cluster of facts, where at least many of the

62. That's not to say that this picture denies that there could be *other* reasons not to eat the fish (e.g., that the fish is over-salted). But it does hold that the fact that the fish contains salmonella enjoys priority relative to the considerations that this picture takes to be in the background of explaining why *it* is a reason (e.g., the fact that salmonella is harmful to humans).

63. Cf. Fogal (2016).

facts among this cluster are roughly on a par, explanatorily speaking. Thus, for example, the fact that fish contains salmonella itself only provides support for refraining from eating it in conjunction with the fact that salmonella is harmful to humans (and perhaps further facts too). To put things another way, the two facts jointly explain why Lois has reason (where ‘reason’ here is a mass noun, as distinct from the count noun ‘reason(s)’) to refrain from eating. Moreover, neither of these facts enjoys privileged status over the other in explaining why Lois has reason to refrain from eating.⁶⁴ This means that it’s a mistake to expect an answer as to which of these two facts is fundamentally “the reason” (where ‘reason’ is a count noun) for Lois to refrain from eating.⁶⁵

The cluster view can nevertheless acknowledge that we often do felicitously *cite* atomic facts as ‘reasons’, without mentioning the other facts that constitute the remainder of the relevant cluster. However, it maintains that which fact from the cluster is most naturally cited as ‘the reason’ does not reflect any deep metaphysical or explanatory primacy, but rather just which fact is most conversationally relevant—a factor that varies across conversational contexts without any change in the underlying normative facts. For instance, if we (the conversational participants) both already know that salmonella is harmful to humans, but you’re not yet aware that the fish contains salmonella, it’s more helpful for me to cite the fact that the fish contains salmonella if you ask me to give you “the reason” for Lois not to eat the fish. But equally, if things are reversed, and you know that the fish contains salmonella but not that salmonella is harmful to humans, it’s more helpful for me to cite the fact that salmonella is harmful as the reason for her to refrain from eating.

Given the cluster view, Lois’s position in the original version of the case—which, recall, is that of knowing that the fish contains salmonella, but not being in a position to know that salmonella is harmful to humans—isn’t importantly different from that of someone (call her Louise) who is in the converse position of knowing that salmonella is harmful to humans, but not being in a position to know that the fish contains salmonella. Neither is in a position to know all the facts that, jointly, provide normative support for refraining from eating. The atomic view takes it that Lois is in a position to know “the reason” to refrain from eating, but Louise isn’t. This is why Lord thinks Louise’s case is dealt with easily

64. It isn’t, say, that the fact F1 that the fish contains salmonella explains why she has reason to refrain, and the fact F2 that salmonella is harmful to human explains *why F1 explains why* she has reason to refrain. Rather, F1 and F2 simply combine to explain why she has reason to refrain, on a par with each other.

65. The cluster view doesn’t necessarily say that there are *never* any “background conditions” that explain why some cluster of facts provide normative support without being part of the cluster itself—compare what Schroeder (2007: Ch. 2) calls the “no background conditions” view—though it is compatible with this stronger view. It does, however, say that at least many facts that those who assume the atomic view treat as background conditions are in fact part of the cluster.

by the epistemic condition, but Lois's isn't, and so requires a separate "practical condition."⁶⁶ But given the cluster view, this is a mistake. Whatever epistemic condition deals with Louise's case should be equally capable of handling Lois's: no practical condition is needed. This is appealing. Intuitively, pointing to Lois's lack of knowledge about how to use a fact as a reason is an unduly roundabout diagnosis of her problem. Fundamentally, her problem is epistemic, not practical: it's that she doesn't know (and isn't in a position to know) that salmonella is harmful to humans.

6.4. *The New Challenge(s)*

It might now seem that, at least if the cluster view is correct, we have solved our problem: we only need an epistemic condition after all. But in fact, as we'll now show, this falls short of a precisification of the claim that substantive rationality consists in responding to evidence-relative reasons.

To begin with, we haven't said how to reformulate the epistemic condition to apply to the cluster view rather than the atomic view. On the atomic view, there's a list of atomic facts that each constitute reasons, and you meet the epistemic condition for the relevant reason just in case you stand in the right epistemic relation (e.g., being in a position to know) to the atomic fact that constitutes it. But how should the epistemic condition be captured on the cluster view?

We might try saying that in order to be (substantively) rationally required to respond to some cluster of facts that provides support for Φ -ing, you have to stand in the right epistemic relation (e.g., being in a position to know) to *each* fact in the cluster. But there is cause to be dissatisfied with this simple proposal, for at least two reasons.

First, if the cluster of facts gets really large, it will be unrealistic to expect you to be in a position to know *everything* that jointly contributes to providing support for Φ -ing in order to be rationally required to Φ . In such cases, you will be rationally required to Φ even though you don't meet the epistemic condition as stated, and so the stated epistemic condition seems too strong.

Second, consider cases where one is (rationally) misled about the *normative* question of whether some fact or cluster of facts supports Φ -ing, where this normative ignorance doesn't rest on more fundamental descriptive ignorance. Arguably, when one is aware of some cluster of facts that as a matter of fact decisively support Φ -ing, but rationally believes that those facts do not decisively support Φ -ing, one need not be irrational in failing to Φ . But the fact that the cluster decisively supports Φ -ing seems to be one fact that can't itself be part of

66. Cf. Lord (2018: 23–27).

the cluster. Thus, you're still in a position to know each fact in the cluster. In such cases, you're not rationally required to Φ even though you do meet the epistemic condition as stated, and so the stated epistemic condition seems too weak.

It may be, then, that there is no way of formulating the epistemic condition in a simple and straightforward way for the cluster view. To evaluate whether you're rationally required to respond to the normative support provided by some cluster of facts, we may have to engage in a more holistic evaluation of your epistemic position—asking whether there's *enough* relevant information amongst your total evidence such that it's reasonable to deem you irrational if you fail to be sensitive to the normative support in question.

Maybe this imprecision isn't a problem. However, there is a deeper sense in which once the cluster view is in place, the ideology of (count-noun) reasons can be seen to have its limitations. For the way that reasons-talk is deployed, at least by philosophers, often tacitly encodes the atomic view of normative support. Take even a simple and paradigmatic example of a reasons-claim, like the claim that the fact that it's raining is a reason for you to fetch your umbrella. If this claim is understood as saying that the atomic fact that it's raining, all on its own, provides normative support for your fetching your umbrella—and this is how philosophers often (tacitly) understand it—then it is highly doubtful. For it seems that it's only in conjunction with other facts—that you need to go outside, that you don't want to get wet (or that something bad would happen if you got wet), and so on—that the fact that it's raining provides normative support for fetching your umbrella.

The complications that we've been exploring for precisifying the slogan that substantive rationality consists in responding to evidence-relative reasons are a symptom of this broader limitation of reasons-talk. If it's assumed that the sources of normative support are (count-noun) reasons, where those are taken paradigmatically to consist in atomic facts, then it's going to seem that as long as one stands in the right epistemic relation to those atomic facts, one also stands in the right epistemic relation to the corresponding strain of normative support (in order to be substantively rationally required to respond to it). But that's what generates counterexamples like *Lois's Fish*.

On the cluster view, by contrast, seeing whether one stands in the right epistemic relation to any particular strain of normative support (in order to be rationally required to respond to it) will require a broader evaluation of one's epistemic position with regard to a whole cluster of facts. This blocks counterexamples like *Lois's Fish*. But it also makes the slogan that rationality consists in responsive to evidence-relative *reasons* somewhat misleading, insofar as reasons are identified with atomic facts.

In light of this, we have at least two options. We could try to stipulate a way of understanding (count-noun) 'reasons' that decouples the philosopher's

deployment of the notion of a reason from the assumption of the atomic view. (For example, we could say that *the reason* for fetching your umbrella is constituted by the whole cluster of facts that provide normative support for doing so.⁶⁷) Alternatively, we could refrain from invoking the ideology of “reasons” in providing a philosophical analysis of (substantive) rationality, and replace it with talk of what we have most *reason* (mass noun) to do. The mass noun ‘reason’ corresponds to what we’ve been calling ‘normative support’, and doesn’t encode the assumption that this normative support is provided by atomic facts, as opposed to clusters of facts.

Which of these ways of talking we adopt is a terminological matter. And whichever way we go, we should still understand the normative support that’s relevant to substantive rationality in a broadly evidence-relative way. Ultimately, then, the challenge we’ve discussed for the view that substantive rationality consists in responding to evidence-relative reasons turns out not so much to be a challenge for the view that it consists in responding to *evidence-relative* reasons, but rather a challenge for the view that it consists in responding to evidence-relative *reasons*—at least where ‘reasons’ are being understood as paradigmatically consisting in atomic facts. Thus, even the best version of the slogan that rationality consists in responding to reasons—namely, that substantive rationality consists in responding to evidence-relative reasons—is at best misleading, and at worst mistaken.

Acknowledgments

For comments or discussion we are grateful to Marc-Kevin Daoust, Alex Gregory, Chris Howard, Zoë Johnson King, Alex King, Stephanie Leary, Olle Risberg, Keshav Singh, Chris Tucker, Daniel Wodak, several anonymous referees (especially two very helpful referees for *Ergo*), and audiences at the 2019 Frankfurt Metaethics Conference, the 2019 Northeast Normativity Workshop, the Uppsala Higher Seminar in Theoretical Philosophy, and the 2020 Central APA.

References

Broome, John (2007). Does Rationality Consist in Responding Correctly to Reasons? *Journal of Moral Philosophy*, 4(3), 349–74.

67. On this way of going, the philosophical notion of a ‘reason’ becomes a semi-technical term, since it’s rare for people to cite the whole cluster of facts that provide normative support as ‘reasons’ in everyday language; they usually just pick a salient *representative* of the cluster.

- Broome, John (2013). *Rationality Through Reasoning*. Wiley-Blackwell.
- Broome, John (in press). Rationality versus Normativity. *Australasian Philosophical Review*. Page references are to the draft available here: <http://users.ox.ac.uk/~sfop0060/pdf/Rationality%20versus%20normativity.pdf>
- Chang, Ruth (2013). Grounding Practical Normativity: Going Hybrid. *Philosophical Studies*, 164(1), 163–87.
- Dancy, Jonathan (2000). *Practical Reality*. Oxford University Press.
- Feldman, Richard (1988). Subjective and Objective Justification in Ethics and Epistemology. *The Monist*, 71(3), 405–19.
- Finlay, Stephen (2014). *Confusion of Tongues: A Theory of Normative Language*. Oxford University Press.
- Fogal, Daniel (2016). Reasons, Reason, and Context. In Errol Lord and Barry Maguire (Eds.), *Weighing Reasons* (74–103). Oxford University Press.
- Fogal, Daniel (2018). On the Scope, Jurisdiction, and Application of Rationality and the Law. *Problema*, 12, 21–57.
- Fogal, Daniel (2020). Rational Requirements and the Primacy of Pressure. *Mind*, 129(516), 1033–70.
- Foot, Philippa (1972). Morality as a System of Hypothetical Imperatives. *Philosophical Review*, 81(3), 305–16.
- Gert, Joshua (2004). *Brute Rationality: Normativity and Human Action*. Cambridge University Press.
- Gibbons, John (2010). Things That Make Things Reasonable. *Philosophy and Phenomenological Research*, 81(2), 335–61.
- Gibbons, John (2013). *The Norm of Belief*. Oxford University Press.
- Henning, Tim (2014). Normative Reasons Contextualism. *Philosophy and Phenomenological Research*, 88(3), 593–624.
- Kelly, Thomas (2007). Evidence and Normativity: Reply to Leite. *Philosophy and Phenomenological Research*, 75(2), 465–74.
- Kiesewetter, Benjamin (2017). *The Normativity of Rationality*. Oxford University Press.
- Kolodny, Niko (2005). Why Be Rational? *Mind*, 114(455), 509–63.
- Kolodny, Niko (2007). How Does Coherence Matter? *Proceedings of the Aristotelian Society*, 107(3), 229–63.
- Lord, Errol (2014). The Coherent and the Rational. *Analytic Philosophy*, 55(2), 151–75.
- Lord, Errol (2018). *The Importance of Being Rational*. Oxford University Press.
- Markovits, Julia (2014). *Moral Reason*. Oxford University Press.
- Neta, Ram (2015). Coherence and Deontology. *Philosophical Perspectives*, 29, 284–304.
- Parfit, Derek (2011). *On What Matters (Vol. 1)*. Oxford University Press.
- Portmore, Douglas W. (2011). *Commonsense Consequentialism: Wherein Morality Meets Rationality*. Oxford University Press.
- Rawls, John (1971). *A Theory of Justice*. Harvard University Press.
- Scanlon, Thomas M. (1998). *What We Owe to Each Other*. Harvard University Press.
- Scanlon, Thomas M. (2007). Structural Irrationality. In Geoffrey Brennan, Robert Goodin, Frank Jackson, and Michael Smith (Eds.), *Common Minds: Themes from the Philosophy of Philip Pettit* (84–103). Oxford University Press.
- Schroeder, Mark (2004). The Scope of Instrumental Reason. *Philosophical Perspectives*, 18, 337–64.
- Schroeder, Mark (2007). *Slaves of the Passions*. Oxford University Press.

- Schroeder, Mark (2009). Means-End Coherence, Stringency, and Subjective Reasons. *Philosophical Studies*, 143(2), 223–48.
- Schroeder, Mark (2011). What Does It Take to ‘Have’ a Reason? In Andrew Reisner and Asbjørn Steglich-Petersen (Eds.), *Reasons for Belief* (201–22). Cambridge University Press.
- Schroeder, Mark (2018). Getting Perspective on Objective Reasons. *Ethics*, 128(2), 289–319.
- Setiya, Kieran (2004). Against Internalism. *Noûs*, 38(2), 266–98.
- Setiya, Kieran (2007). Cognitivism about Instrumental Reason. *Ethics*, 117(4), 649–73.
- Sylvan, Kurt (2015). What Apparent Reasons Appear to Be. *Philosophical Studies*, 172(3), 587–606.
- Sylvan, Kurt (in press). Respect and the Reality of Apparent Reasons. *Philosophical Studies*. Published Online Early View; doi: 10.1007/s11098-020-01573-1. Advance online publication. <https://doi.org/10.1007/s11098-020-01573-1>
- Vogelstein, Eric (2012). Subjective Reasons. *Ethical Theory & Moral Practice*, 15(2), 239–57.
- Way, Jonathan (2009). Two Accounts of the Normativity of Rationality. *Journal of Ethics and Social Philosophy*, 4(1), 1–8.
- Way, Jonathan (2012). Explaining the Instrumental Principle. *Australasian Journal of Philosophy* 90(3), 487–506.
- Wedgwood, Ralph (2017). *The Value of Rationality*. Oxford University Press.
- Whiting, Daniel (2014). Keep Things in Perspective: Reasons, Rationality and the A Priori. *Journal of Ethics & Social Philosophy*, 8(1), 1–22.
- Williams, Bernard (1981). Internal and External Reasons. In Bernard Williams, *Moral Luck* (101–13). Cambridge University Press.
- Wodak, Daniel (2019). An Objectivist’s Guide to Subjective Reasons. *Res Philosophica*, 96(2), 229–44.
- Worsnip, Alex (2018). Reasons, Rationality, Reasoning: How Much Pulling-Apart? *Problema*, 12, 59–93.
- Worsnip, Alex (in press). *Fitting Things Together: Coherence and the Demands of Structural Rationality*. Oxford University Press.