

The Unity of Evidence and Coherence

Draft chapter for *Epistemic Dilemmas*,
Edited by Nick Hughes (Oxford University Press, Forthcoming).

An emerging theme in normative theory, including both ethics and epistemology, is that there is an important distinction to be drawn between *substantive* and *structural* requirements of rationality (Scanlon 2007; Worsnip 2018a). The goal of this chapter is to examine this distinction as it applies within the epistemic domain. What is the relationship between substantive and structural requirements of epistemic rationality?

Substantive rationality is a matter of responding to reasons. Scanlon defines a *reason* for something as “a consideration that counts in favor of it” (1998: 17). Given the flexibility of the English word ‘reason’, however, there is no single sense in which reasons count in favor of belief or action (Schroeder 2008). For current purposes, a reason to φ counts in favor of φ -ing in the sense that it contributes towards making it substantively rational to φ . According to *evidentialism* in epistemology, all reasons for belief are constituted by evidence. Given this assumption, substantive rationality in the epistemic domain is a matter of proportioning your beliefs to your evidence.

Structural rationality, in contrast, is a matter of having beliefs and other attitudes that cohere or “fit together” in the right way. Although there is some controversy about which forms of coherence are required for structural rationality, the following requirements are widely endorsed. First, your beliefs should be *logically coherent* in the sense that they are logically consistent and closed under logical consequence. Second, your degrees of belief or “credences” should be *probabilistically coherent* in the sense that they conform to the axioms of the probability calculus. And third, your beliefs should be *meta-coherent* in the sense that your object-level beliefs cohere with your meta-level beliefs about which beliefs you should hold.

It is plausible that both substantive and structural requirements are built into our ordinary conception of epistemic rationality. An epistemically rational agent is one whose beliefs are both coherent and proportioned to her evidence. Anyone whose beliefs are either incoherent or unsupported by her evidence has thereby violated some requirement of epistemic rationality.

Some theories of epistemic rationality impose no substantive requirement to respect your evidence, but these theories have little to recommend them. On pure versions of coherentism, for example, epistemic rationality is simply a matter of having beliefs that cohere with each other. Similarly, according to subjective Bayesianism, there is nothing more to epistemic rationality than having credences that are probabilistically coherent and updated by Bayesian conditionalization. These views are subject to compelling counterexamples. Consider Magic Feldman, who combines Magic Johnson’s beliefs with Richard Feldman’s experiences

(Feldman 2003: 68). Although his beliefs are coherent, they are not epistemically rational. Epistemic rationality requires that your beliefs cohere not only with each other, but also with the evidence provided by your sensory experience.

Other theories of epistemic rationality abandon any structural requirements of coherence. On a global version of phenomenal conservatism, for example, epistemic rationality is simply a matter of believing whatever seems true (Huemer 2001). Since there are no structural constraints on which propositions can seem true, this view implies that epistemically rational agents can hold beliefs that are wildly incoherent (Smithies 2019: Ch. 12). Arguably, this compromises the value of epistemic rationality by failing to preserve the structural features that are distinctive of epistemically rational agents. Any conception of epistemic rationality that deserves the name must recognize at least some structural requirements of coherence.

This chapter begins from the plausible assumption that epistemic rationality requires both coherence and respecting your evidence. The main goal of the chapter is to examine the relationship between these requirements. What is the relationship between the substantive requirement to respect your evidence and the structural requirement to be coherent?

Let's define *bifurcationism* as the thesis that the substantive and structural requirements of epistemic rationality are distinct and *sui generis* in the sense that neither can be explained as a consequence of the other. On this view, there are two fundamentally different kinds of requirements of epistemic rationality, neither of which can be reduced to the other. One argument for bifurcationism is that these requirements can come into conflict when your evidence supports incoherent beliefs. Thus, Alex Worsnip (2018a) argues that your evidence supports meta-incoherent beliefs when it is misleading about itself. Such cases threaten to generate *epistemic dilemmas* in which you are guaranteed to violate either the substantive requirement to respect your evidence or the structural requirement of meta-coherence.

The central thesis of this chapter is that there can be no conflict between evidence and coherence. This is because there are structural constraints on the evidential support relation, which guarantee that your evidence never supports incoherent beliefs. Any epistemically rational agent who respects her evidence is thereby guaranteed to be coherent, since the evidence always supports coherent beliefs. According to this version of *unificationism*, there is fundamentally just one requirement of epistemic rationality, which incorporates both substantive and structural dimensions. This is the evidentialist requirement to hold beliefs that cohere with substantive facts about your evidence in accordance with structural facts about the evidential support relation.

The challenge remains to explain where Worsnip's argument goes wrong. I argue that your evidence cannot be misleading about itself because the facts about what your evidence is and what it supports are always *self-evident* in the sense that they are certain given your evidence. On this view, respecting your evidence guarantees that your beliefs are not only coherent, but also meta-coherent.

Moreover, I explain away intuitions to the contrary by appealing to a distinction between ideal and non-ideal requirements of epistemic rationality.

Here is the plan. §1 presents a puzzle that arises from the bifurcationist assumption that there can be conflicts between evidence and coherence. §2 critiques various bifurcationist solutions to the puzzle, while §3 raises more general problems with the bifurcationist assumption that generates the puzzle in the first place. §4 presents my unificationist solution to the puzzle, while §5 defends it by explaining away the intuitions that seem to favor bifurcationism. Finally, §6 concludes by explaining how my unificationist solution avoids positing epistemic dilemmas.

1. A Puzzle about Evidence and Coherence

This section presents a puzzle that arises from the bifurcationist assumption that there can be conflicts between evidence and coherence. The puzzle is generated by cases in which the substantive requirement to respect your evidence appears to conflict with the structural requirement to be coherent because your evidence supports incoherent beliefs. Such cases threaten to yield epistemic dilemmas in which epistemic rationality issues logically inconsistent requirements, but this is hard to stomach. Is it really tolerable to suppose that epistemic rationality sometimes requires us to do what is logically impossible?

The puzzle can be presented in the form of a paradox:

- (1) Your evidence sometimes supports only incoherent beliefs.
- (2) Epistemic rationality always requires that your beliefs are supported by your evidence.
- (3) Epistemic rationality always requires that your beliefs are coherent.
- (4) Epistemic rationality never issues inconsistent requirements.

Each of these four claims is individually plausible but they are jointly inconsistent. If your evidence supports only incoherent beliefs, then you're guaranteed to violate a requirement of epistemic rationality, since you cannot respect your evidence while also remaining coherent. In such cases, you face an epistemic dilemma in which epistemic rationality issues inconsistent requirements.

Worsnip (2018a) explores a version of this puzzle that is generated by *higher-order evidence*: that is, evidence about your own evidence. The key idea is that when you have misleading higher-order evidence, your total evidence can be misleading about itself. For example, your total evidence can support the proposition that p while also supporting the higher-order proposition that your total evidence doesn't support p . In that case, your evidence makes it substantively rational to believe the conjuncts of an abominable conjunction:

p but I shouldn't believe p because my total evidence doesn't support p .

This generates a version of Moore's paradox, since the abominable conjunction is consistent and yet it nevertheless seems epistemically irrational to believe it (Smithies 2012; Horowitz 2014). Someone who believes the conjuncts of this conjunction is *epistemically akratic* in a way that seems irrational. Just as it seems irrational to act in conflict with your beliefs about how you should act, so it seems irrational to hold beliefs that conflict with your beliefs about what you should believe. Akrasia, whether practical or epistemic, is a paradigmatic form of structural irrationality. Nevertheless, Worsnip argues, your evidence makes it substantively rational to be epistemically akratic. Hence, there is an apparent conflict between structural and substantive requirements of epistemic rationality.

Worsnip illustrates the point with the following example:

Miss Marple and Mabel. Miss Marple is a detective who is famously good at assessing evidence. Miss Marple is investigating a murder that took place at the mansion on the hill, and she takes her great niece Mabel along with her. Miss Marple and Mabel set about the mansion collecting clues. Unfortunately, in their initial sweep of the house, nothing that they learn offers any kind of significant support to any particular hypothesis about who committed the crime. As part of her training of Mabel as her apprentice, after they have finished examining a crime scene, Miss Marple always tells Mabel what [is] her own assessment of what the evidence supports. On this occasion, Miss Marple makes an uncharacteristic error, and declares to Mabel, "the clues lying around the house that you have seen up to this point support believing that the vicar did it". (2018a: 24)

Although the clues themselves are not probative, Miss Marple's expert testimony provides higher-order evidence that the clues incriminate the vicar. According to Worsnip, Mabel's total evidence supports agnosticism about whether the vicar is guilty, while also supporting the higher-order belief that her evidence supports believing this conclusion. And yet it seems structurally irrational for Mabel to hold meta-incoherent beliefs, which she might express as follows: "Given my evidence, I should believe the vicar is guilty, but I refuse to believe this." Hence, Mabel seems faced with an epistemic dilemma, since she cannot respect her evidence while also maintaining meta-coherence.

David Christensen (2007) gives similar examples in which you have misleading evidence about your own logical reasoning abilities:

Reason-Distorting Drugs. Suppose that I work out my proof of T after having coffee with my friend Jocko. Palms sweaty with the excitement of logical progress, I check my work several times, and decide that the proof is good. But then a trusted colleague walks in and tells me that Jocko has been surreptitiously slipping a reason-distorting drug into people's coffee – a drug whose effects include a strong propensity to reasoning errors in 99% of those who have been dosed (1% of the population happen to be immune).

He tells me that those who have been impaired do not notice any difficulties with their own cognition – they just make mistakes; indeed, the only change most of them notice is unusually sweaty palms. (2007: 10)

Christensen argues that it's irrational to be certain that T while doubting the reasoning that leads to this conclusion, since this is a form of meta-incoherence. However, your evidence that you have ingested the drug makes it rational to doubt your own reasoning. And yet probabilistic coherence requires being certain that T, since T is a logical truth that has probability 1. This means that you are guaranteed to violate one of the following requirements of epistemic rationality:

LOGIC: An agent's beliefs must respect logic by satisfying (some version of) probabilistic coherence.

EVIDENCE: An agent's beliefs (at least about logically contingent matters) must be proportioned to the agent's evidence.

INTEGRATION: An agent's object-level beliefs must reflect the agent's meta-level beliefs about the reliability of the cognitive processes underlying her object-level beliefs. (2007: 20)

You cannot respect your higher-order evidence about your own reasoning, while also integrating your first-order and higher-order beliefs in a meta-coherent way, without thereby violating probabilistic coherence. Hence, you seem faced with an epistemic dilemma in which you are guaranteed to violate a requirement of epistemic rationality.

Some epistemologists deny that meta-coherence is a genuine requirement of epistemic rationality (Coates 2012; Weatherson 2019; Lasonen-Aarnio 2020). And yet this is not enough to dissolve the puzzle. We can generate puzzle cases through apparent conflicts between LOGIC and EVIDENCE without invoking INTEGRATION. In the simplest cases, you receive expert testimony that T is false, when in fact T is a logical truth. In such cases, you need not have any misleading higher-order evidence about your own reasoning abilities. Even so, it seems plausible that the expert testimony gives you misleading evidence about logic. But you cannot respect this evidence without violating probabilistic coherence.

Some epistemologists deny that probabilistic coherence is a genuine requirement of epistemic rationality (Weatherson 2019). But this is still not enough to dissolve the puzzle, since we can generate apparent conflicts between evidence and coherence without making any assumptions about the exact form of these coherence requirements. Whichever beliefs are prohibited by these coherence requirements, we can imagine receiving expert testimony that provides misleading evidence that those beliefs are true. Hence, we cannot dissolve the puzzle by rejecting specific coherence requirements unless we make the implausible claim there are no coherence requirements at all.

2. Bifurcationist Solutions

This section examines various bifurcationist strategies for solving the puzzle. One of the main arguments for bifurcationism is that the structural and substantive requirements of epistemic rationality are distinct because they conflict when your evidence supports incoherent beliefs. Hence, proponents of bifurcationism typically accept claim (1) and try to solve the puzzle in some other way. This section examines four different bifurcationist strategies for solving the puzzle and finds problems with them all.

2.1. Epistemic Dilemmas

The first strategy is to reject claim (4) that epistemic rationality never issues inconsistent requirements (Hughes 2019). On this view, epistemic rationality always requires respecting your evidence, and always requires coherence, although you cannot always satisfy both requirements. When your evidence supports incoherent beliefs, you ought to hold these beliefs because they are supported by your evidence, but at the same time you ought not to hold them because they are incoherent. These cases are *dilemmas* in the strict sense that they instantiate the formal schema below:

Dilemmas: $O\varphi$ & $O\neg\varphi$

There is no logical contradiction in the claim that there can be dilemmas, since $O\neg\varphi$ doesn't imply $\neg O\varphi$. Nevertheless, this claim yields logical contradiction when combined with standard principles in deontic logic, such as the following:

Ought-Implies-May: $O\varphi \rightarrow P\varphi$

Duality: $P\varphi \rightarrow \neg O\neg\varphi$

Together, these two principles imply that $O\varphi \rightarrow \neg O\neg\varphi$, which is inconsistent with the thesis that there are dilemmas.

Similarly, we can derive a logical contradiction from the claim that there are dilemmas given the three principles below:

Agglomeration: $O\varphi$ & $O\psi \rightarrow O(\varphi \& \psi)$

Ought-Implies-Can: $O\varphi \rightarrow C\varphi$

Can't do the Impossible: $\neg C(\varphi \& \neg\varphi)$

Suppose there are strict dilemmas in which I ought to φ and I ought not to φ . Given the agglomeration principle, I ought both to φ and to refrain from φ -ing. Given the ought-implies-can principle, it follows that I can both φ and refrain from φ -ing. And yet I cannot both φ and refrain from φ -ing, since this is logically impossible. Once again, we derive a logical contradiction.

Hughes (2019) defends the coherence of dilemmas by rejecting the agglomeration principle and the ought-implies-may principle. And yet rejecting these principles comes with a significant theoretical cost in simplicity or explanatory power. Either we cannot explain the correctness of patterns of deontic reasoning that are validated by these principles or else we must complicate the explanation by restricting these principles so they don't apply in full generality. I doubt these costs are worth paying to make sense of epistemic dilemmas.

Moreover, there is a problem with the very idea of an epistemic dilemma in which epistemic rationality issues inconsistent requirements. Epistemic rationality is a dimension of epistemic value: there is always some value in being epistemically rational, though it may be outweighed by conflicting values. To be epistemically rational is to satisfy all the requirements of epistemic rationality. But if there are strict dilemmas in which epistemic rationality issues inconsistent requirements, then it is logically impossible to satisfy all of them at once. And there can be no value in satisfying all the requirements of epistemic rationality when it is logically impossible to do so, since there can be nothing of value in a logically impossible scenario. Hence, the very idea of an epistemic dilemma compromises the value of epistemic rationality.

Proponents of epistemic dilemmas may retreat to the claim that epistemic rationality is valuable only insofar as it is logically possible to achieve it. Rather than compromising the value of epistemic rationality, however, it seems preferable to maintain that epistemic rationality is always valuable by denying that it ever issues logically inconsistent requirements. Arguably, it is an adequacy constraint on any theory of epistemic rationality that it should delineate something that always has epistemic value. This is enough to motivate the search for a solution to our puzzle that avoids positing epistemic dilemmas.

2.2. *Prima Facie* Requirements

The second option is to treat the structural and substantive requirements of epistemic rationality as mere *prima facie* requirements to maximize various rational ideals, which can come into conflict (Christensen 2007). On this view, there can be conflicts between the rational ideals of coherence and respecting your evidence. Other things equal, you should maximize each of these rational ideals, but things are unequal when they conflict. In such cases, you're not required to do the impossible: namely, to be perfectly coherent while also perfectly respecting the evidence. Rather, you are required to find the best overall balance between these rational ideals. As Christensen writes, "We're quite familiar with other ideals that operate as values to be maximized, yet whose maximization must in certain cases be balanced against, or otherwise constrained by, other values" (2007: 24).

On this view, conflicts between rational ideals do not generate strict dilemmas in which you're required both to believe and to refrain from believing one and the same proposition. At best, these conflicts generate dilemmas in the colloquial sense: that is, hard cases in which it's not obvious how you should weigh conflicting values against each other. And yet some conflicts may generate easy

cases in which it seems clear how they should be resolved. For example, Christensen claims that it's not rationally optimal to be probabilistically coherent in the face of higher-order evidence of your own cognitive imperfection. Instead, you should reduce your confidence in logical truths when you have higher-order evidence that your logical reasoning is impaired – say, by reason-distorting drugs. In such cases, the value of probabilistic coherence is outweighed by the value of respecting the evidence about your cognitive abilities.

How does this view solve the puzzle? On this view, (2) and (3) are false because epistemic rationality doesn't always require respecting your evidence and also remaining coherent. When your evidence supports incoherent beliefs, it's permissible to weigh these competing ideals against each other. Epistemic rationality requires only that your beliefs reflect some optimal balance between coherence and respecting the evidence.

This view gives a more plausible treatment of hard cases than treating them as strict dilemmas. Consider Sartre's (1946) example of the student deciding whether to join the French resistance to the German occupation or instead to stay at home with his mother who depends on him. It seems absurd to suggest that he should do both things when this is clearly impossible. It is vastly more plausible that he should pursue whichever course of action is supported by stronger reasons, although it may be hard to ascertain where the balance of reasons lies. If the reasons are equally balanced, of course, then it may be permissible to pursue either course of action.

Unfortunately, we cannot extend the same treatment to the examples that concern us. In Sartre's example, there is a conflict between substantive reasons for action, which derive from competing considerations about the value of helping your family versus the value of defending your nation. What you ought to do in such cases depends on the overall balance of substantive reasons. Our examples, in contrast, do not concern conflicts between substantive reasons, but between the substantive requirement to respect your reasons for belief and the structural requirement to be coherent. We cannot resolve this conflict by appealing to the overall balance of substantive reasons.

The objection can be posed as a dilemma. Do we have substantive reasons to be coherent or not? If not, we have no model for weighing substantive against structural considerations in determining what you ought to believe. This imposes some theoretical pressure to say that we have substantive reasons to be coherent after all. But if so, the substantive reasons for belief provided by your evidence must be weighed against the substantive reasons for belief provided by structural considerations about coherence. And this is to reject the evidentialist assumption that all substantive reasons for belief are provided by evidence alone. Evidentialism is not sacrosanct, of course, but it is plausible enough to motivate the search for an alternative solution.

2.3. Rational Indeterminacy

The third option is that it's *indeterminate* what you should believe when substantive and structural requirements of epistemic rationality come into conflict (Leonard 2020). On this view, conflicts between evidence and coherence do not generate epistemic dilemmas in which it's determinately true that you ought and ought not to believe one and the same proposition. On the contrary, these are cases in which it's indeterminate what you ought to believe.

In a supervaluationist framework, it's determinately true that you ought to hold a belief just in case the belief is required by every maximally consistent way of resolving the requirements of epistemic rationality. When your evidence supports incoherent beliefs, it's indeterminate what you ought to believe because there are multiple ways of resolving the requirements of epistemic rationality, which diverge in what they require. Some resolutions require coherence, while others require respecting your evidence, although none require both. Hence, there are no epistemic dilemmas, since it is determinately false that you ought both to respect your evidence and to be coherent in such cases.

How does this view solve our puzzle? On this view, neither (2) nor (3) is determinately true, since it's not determinately true that epistemic rationality requires you to respect your evidence, or that it requires you to be coherent, when your evidence supports incoherent beliefs. On the contrary, it's indeterminate what epistemic rationality requires in such cases. Indeed, it's determinately false that epistemic rationality requires both coherence and respecting your evidence in such cases. Hence, the conjunction of (2) and (3) is determinately false.

I have no objection to the general claim that it's sometimes indeterminate what you should believe or do. I see no special reason to suppose that ethics and epistemology are immune from indeterminacy. Arguably, though, this solution countenances too much normative indeterminacy, since it implies that the normative facts are indeterminate whenever conflicts arise between substantive and structural requirements of rationality. Hence, this solution cannot vindicate intuitive verdicts about what we should believe and do in such cases.

Consider a pilot who has misleading evidence that she is suffering from hypoxia, but who calculates correctly that she has enough fuel to take a scenic detour en route to her final destination (Elga 2013). Although her evidence supports this conclusion, it seems reckless for her to decide on this basis to take the detour, rather than flying directly to her destination, given the higher-order evidence that she is cognitively impaired by hypoxia. Intuitively, she shouldn't believe or act on the conclusion that is supported by her evidence. On the current proposal, however, it is indeterminate whether she should disrespect her evidence in this way or instead violate meta-coherence by believing a conclusion that she regards as probably based on mistaken reasoning. Many will regard this as a counterintuitive prediction of the theory.

Another problem is that this solution cannot vindicate the plausible theoretical principle that epistemic rationality always requires respecting your evidence and remaining coherent. Bifurcationism promises to accommodate the

normative force of both requirements by allowing that they can conflict. And yet this solution cannot maintain that both requirements retain their normative force when they conflict, since it implies that it's indeterminate which one you should comply with. Worse, it implies that it's determinately true that you shouldn't comply with both requirements when they conflict. In other words, it's determinately false that you should always respect your evidence while remaining coherent. This is implausible enough to motivate the search for a solution to our puzzle that can accommodate the normative force of both requirements.

2.4. Equivocation

The fourth option is to deny that there is any single sense in which you ought to respect your evidence and to remain coherent when your evidence supports incoherent beliefs (Worsnip 2018a). On this view, these cases involve conflicts between normative domains, rather than within a single normative domain. We cannot understand the distinction between substantive and structural requirements in terms of a single normative concept of epistemic rationality that governs what we should believe. Instead, our beliefs are governed by two distinct and fundamentally different kinds of normative requirements that cannot be stated using the same normative concept.

How does this view solve our puzzle? On this view, the appearance of a puzzle arises from an equivocation between two distinct normative concepts. There is no single normative concept of epistemic rationality in terms of which we are always required both to respect our evidence and to be coherent. Hence, there is no single interpretation on which both (2) and (3) are true.

This view has several advantages. First, it avoids strict dilemmas, which removes any pressure to revise deontic logic. Second, by distinguishing normative requirements, it promises to vindicate the plausible theoretical principle that you are always required to respect your evidence and to remain coherent. Third, since these requirements are fundamentally distinct, there is no danger of compromising evidentialism as a thesis about substantive reasons for belief by encroachment from non-evidential considerations about coherence. And, finally, since these requirements are incommensurable, there is no commitment to any common scale on which they can be weighed against each other.

Despite these attractions, however, problems remain. One problem is that it's hard to make sense of intuitions about what you should do in conflict cases. According to the equivocation strategy, there is one sense in which you ought to respect your evidence and another sense in which you ought to remain coherent. But this doesn't capture the intuitive sense that it's *better* for the pilot to maintain meta-coherence, rather than respecting her evidence, when she acquires the higher-order evidence that she is hypoxic.

Indeed, the equivocation strategy doesn't even provide us with the conceptual resources for asking whether it's better to respect evidence or maintain coherence in conflict cases. After all, we have no single normative concept in terms of which these conflicting requirements can be weighed against each other. We

can ask what substantive rationality requires, and we can ask what structural rationality requires, but we cannot coherently ask which of these requirements are better to follow. And yet it does seem like we can coherently ask this question, or something like it, when we consider the hypoxia case.

Perhaps the deepest problem is that we lose the attractive idea that there is any unified virtue of epistemic rationality that requires both coherence and respecting your evidence. According to the equivocation strategy, there is a substantive requirement to respect your evidence, and a structural requirement to be coherent, but there is no single concept of epistemic rationality that unifies these requirements. Moreover, we cannot reconstruct a unified virtue of epistemic rationality by simply conjoining these requirements. After all, there is always some value in epistemic rationality, but there is no value in satisfying both substantive and structural requirements when they conflict, since this is logically impossible. There is no value in what is logically impossible.

3. Problems for Bifurcationism

The previous section raised problems for bifurcationist attempts to solve the puzzle, whereas this section raises more general problems for the bifurcationist assumption that generates the puzzle in the first place. I argue that we should reject the bifurcationist assumption that your evidence can support incoherent beliefs. Instead, we should prefer a more unified conception of epistemic rationality, which builds the structural requirements of coherence into the substantive requirement to respect your evidence. We should prefer unificationism to bifurcationism because (i) it is more parsimonious, (ii) it better explains the value of coherence, and (iii) the distinction between substantive and structural requirements is dubiously intelligible in the first place. Let's take these three points in reverse order.

3.1. No Intelligible Distinction

My first argument is that there is no intelligible distinction between substantive and structural requirements of epistemic rationality. We cannot ultimately make sense of the idea that some requirements of epistemic rationality are purely substantive, rather than structural, or vice versa. As Ralph Wedgwood writes, this is "a distinction without a difference" (2017: 11).

Consider first the requirement to proportion your beliefs to your evidence. Can we make any sense of the idea that this requirement is purely substantive rather than structural? The requirement is that your degree of belief in a proposition should be proportionate with the degree to which your evidence supports that proposition. But the degree to which your evidence supports a proposition depends not only on substantive facts about what your evidence is, but also on structural facts about the evidential support relation. Without invoking the structure of the evidential support relation, we cannot explain how your evidence supports any given proposition to any given degree. Hence, there is a structural dimension built into the requirement to proportion your beliefs to the evidence.

Why is it, for example, that your evidence never supports contradictions? This is because there are logical constraints built into the structure of the evidential support relation. Your evidence supports a contradiction, p and not- p , only if it supports both conjuncts. But your evidence cannot support a proposition while also supporting its negation. These logical constraints on the evidential support relation are best captured within a probabilistic framework, according to which your evidence supports a proposition only if it is more probable than not on your evidence that it is true. It cannot be that a proposition and its negation are both more probable than not to be true, since the probability of the disjunction must sum to one. Hence, your evidence never supports contradictory propositions.

Now consider the requirement to be coherent. Can we make any sense of the idea that this requirement is purely structural rather than substantive? Coherence requires that your beliefs stand in certain relations to your other beliefs and mental states. This requirement has a structural dimension, which concerns the structure of the relations that must hold between your beliefs and other mental states. But it also has a substantive dimension, which concerns the mental states that fall within its scope.

Which mental states fall within the scope of the coherence requirement? The most plausible answer is that epistemic rationality requires your beliefs to cohere not just with your other beliefs, but also with all the other mental states that provide you with evidence. The assumption here is not that your evidence is exhausted by facts about your mental states, but merely that you possess evidence in virtue of being in certain mental states. I claim that your beliefs are fully coherent when they fit together in the right way with all the mental states in virtue of which you possess evidence.

Epistemic rationality requires that your beliefs cohere with each other. This is a consequence of the fact that your beliefs provide you with defeasible evidence, which may be defeated by the evidence provided by other beliefs. Plausibly, however, your experiences provide you with evidence as well as your beliefs. This means that your beliefs can cohere with each other without cohering in all the ways that matter for epistemic rationality. Magic Feldman is a case in point: he is epistemically irrational because his beliefs cohere with each other but not with his experiences. Epistemic rationality requires that his beliefs cohere with all the mental states that provide him with evidence, including his experiences as well as his other beliefs.

Epistemic rationality doesn't require that your beliefs cohere with all your mental states. For example, it is not epistemically irrational to believe a proposition while subdoxastically representing its negation in a mental module. This is because your subdoxastic mental representations, unlike your beliefs, do not provide you with evidence. Epistemic rationality requires that your beliefs cohere with all and only those mental states that provide you with evidence.

How exactly should your beliefs cohere with the mental states that provide you with evidence? The most plausible answer is that the structure of the coherence requirement derives from the structure of the evidential support

relation. Your beliefs should cohere with the mental states that provide you with evidence in accordance with the structure of the evidential support relation. Hence, the coherence requirement for epistemic rationality is none other than the requirement to hold beliefs that cohere with your evidence.

The upshot is that there is no intelligible distinction between substantive and structural requirements of epistemic rationality. Fundamentally, there is just one requirement of epistemic rationality, which incorporates both substantive and structural dimensions. This is the evidentialist requirement to proportion your beliefs to the evidence in the sense that they cohere with substantive facts about your evidence in accordance with structural facts about the evidential support relation. If your beliefs are fully proportioned to the evidence, then they are fully coherent, and also vice versa. Coherence and respecting the evidence are two sides of the same coin.

3.2. The Value of Coherence

My second argument concerns the value of coherence. If unificationism is true, then the value of coherence consists in the value of respecting your evidence. If bifurcationism is true, however, the value of coherence is much harder to explain. What is the value of coherence when it comes at the cost of respecting your evidence? Does it have any genuine epistemic value or does it merely reflect some fetish for neat and tidy belief systems? And if coherence has no genuine epistemic value, then what is the normative force of the structural requirement to be coherent? Do we have any good reason to be coherent?

One answer is that coherence has intrinsic value. This has some plausibility given the unificationist view that the value of coherence consists in the value of respecting your evidence. And yet it is much less plausible that coherence has intrinsic value when it results from disrespecting your evidence. As Niko Kolodny writes, “It seems outlandish that the kind of psychic tidiness that . . . formal coherence enjoins should be set alongside such final ends as pleasure, friendship, and knowledge” (2007: 241).

A second answer is that coherence has instrumental value because it is a means to an end that has intrinsic value – namely, respecting your evidence. Again, this claim has some plausibility given the unificationist view that your evidence always supports coherent beliefs, since your beliefs must be coherent to respect your evidence. And yet it is much less plausible given the bifurcationist view that your evidence can support incoherent beliefs. How is coherence a means to the end of respecting your evidence in such cases? If conflict cases are rare enough, then perhaps coherence is a reliable though not infallible way of respecting your evidence. Once we divorce evidence and coherence, however, it is far from clear that conflict cases will be the exception rather than the rule.

A third answer is that coherence has instrumental value because it is a necessary condition for *agency*, the capacity to act on beliefs and desires. If there are coherence constraints built into the nature of agency, then perhaps we can derive the value of coherence from the value of agency. The assumption is

questionable, since it is not clear why a Lewisian madman cannot act on the basis of beliefs and desires that are completely incoherent (Smithies *et al.* forthcoming). Even granting this assumption, however, problems remain.

One problem is that any coherence constraints on agency must be extremely weak. Perfect coherence cannot be required for agency, since human agents fall well short of this demanding threshold. At best, agency requires some minimal degree of coherence. So what explains the added value of increasing your degree of coherence beyond this minimal threshold? Presumably, it is not necessary for being an agent that you increase your degree of coherence so long as you meet the minimal threshold to qualify as an agent in the first place.

Another general problem with this strategy is that it presupposes the value of agency without explaining it. Why are we entitled to this assumption? As David Enoch (2006) articulates the question, why should we care about *agency*, rather than *schmagency*? As I'll explain, this challenge is especially urgent for proponents of bifurcationism who build coherence constraints into the nature of agency.

Consider Worsnip's (2018a) bifurcationist thesis that your total evidence can support an incoherent set of beliefs. Now combine this with Worsnip's (2018b) claim that a set of beliefs is *incoherent* just in case it is partially constitutive of the nature of belief that any agent is disposed to revise those beliefs under conditions of full self-knowledge. A consequence of these two claims is that there cannot be an ideal epistemic agent whose beliefs are always perfectly proportioned to her evidence. Whenever her evidence supports incoherent beliefs, such an agent holds incoherent beliefs under conditions of full self-knowledge with no disposition to abandon them. But Worsnip's account of incoherence excludes this possibility, since it is inconsistent with the coherence constraints on the nature of belief.

This is a surprising result. It is often thought to be possible in principle, if not in practice, that there could be an ideal epistemic agent whose beliefs are always perfectly proportioned to her evidence. And yet this possibility is excluded by combining Worsnip's bifurcationism with his account of incoherence. More importantly, anyone who bites this bullet faces an awkward normative question. What is the normative significance of the fact that the constitutive nature of belief precludes incoherence under conditions of full self-knowledge? Is it a good thing because it provides us with some protection against incoherence? Or is it a bad thing because it imposes an obstacle that prevents us from proportioning our proportioning our beliefs to the evidence?

The second answer is hard to avoid. The nature of belief, desire, and agency is normatively defective insofar as it excludes the possibility of proportioning your beliefs to the evidence in conflict cases. Given the value of substantive rationality, it would be better not to be an agent with beliefs and desires, since the coherence constraints on agency pose an obstacle to substantive rationality. It would be better to have belief-like and desire-like states that are not subject to these coherence constraints. In Enoch's terms, it would be better to be a "schmagent" with "schmeliefs" and "schmesires", rather than an agent with beliefs and desires.

I conclude that proponents of bifurcationism cannot derive the value of coherence from the value of agency. More generally, I suspect that we cannot explain the value of coherence without endorsing the unificationist view that it reduces to the value of respecting your evidence.

3.3. Occam's Razor

My third and final argument is an appeal to theoretical parsimony. Bifurcationism says that the substantive and structural requirements of epistemic rationality are distinct and *sui generis*. My version of unificationism, in contrast, says there is just one requirement of epistemic rationality, which incorporates both structural and substantive dimensions – namely, to hold beliefs that cohere with your evidence. Hence, parsimony favors this view by an application of Occam's razor: don't multiply requirements of epistemic rationality beyond necessity!

This argument is not conclusive, of course, since Occam's razor permits multiplying requirements of epistemic rationality when it is necessary to do so. Nevertheless, it is enough to impose an argumentative burden on proponents of bifurcationism. Do we have any good reason to divorce the structural requirement of coherence from the substantive requirement to respect your evidence? One argument is that we need to recognize *sui generis* coherence requirements to explain the normative difference between subjects who differ in coherence without respecting their evidence. In response, however, I'll argue that we can explain the intuitive data without appealing to *sui generis* coherence requirements.

Consider three detectives working on a case who disagree about the cause of the victim's death:

- Amy is agnostic about whether it is murder or suicide.
- Beth believes it is murder, rather than suicide.
- Carl believes it is murder, and also believes it is suicide, although he knows it cannot be both murder and suicide.

Let's assume that only Amy succeeds in proportioning her beliefs to the evidence, since the evidence is neutral between murder and suicide. Hence, Beth and Carl fare equally poorly in responding to the evidence. Nevertheless, there is an intuitive sense in which Beth is doing epistemically better than Carl, since at least her beliefs are coherent. Bifurcationism can explain this easily. Although they both violate the substantive requirement to respect their evidence, only Beth satisfies the structural requirement to have coherent beliefs. The challenge for unificationism is to explain the intuitive sense that Beth does epistemically better than Carl without divorcing the structural requirements of coherence from the substantive requirement to respect your evidence.

My response is that Beth is more *reasonable* than Carl in the sense that she manifests better reasoning dispositions, which dispose her to succeed in respecting her evidence in other cases (Lasonen-Aarnio 2010). To see the point, suppose this is a hard case: although the evidence is neutral between murder and suicide, it is

easily confused with a simpler case in which the evidence clearly supports murder, rather than suicide. Beth is disposed to proportion her beliefs to the evidence in the easy case, but not the hard case, whereas Carl is disposed to proportion his beliefs to the evidence in neither case. In that sense, Beth's reasoning dispositions are more responsive to evidence than Carl's. This explains the intuitive sense that Beth does epistemically better than Carl without divorcing requirements of coherence from the substantive requirement to respect the evidence.

I'm assuming that Beth's reasoning dispositions are imperfectly sensitive to evidence in such a way that she respects her evidence in easy cases but not hard cases. Of course, we can stipulate a case in which Beth has coherent beliefs that do not result from evidence-sensitive dispositions at all. But now I lose my intuitive sense that she is doing epistemically better than Carl. Perhaps it's just a lucky coincidence that her beliefs are coherent or perhaps there is some explanation that involves no evidence-sensitive dispositions. Either way, I doubt that the coherence in her belief system reflects anything of epistemic value.

The key challenge for proponents of bifurcationism is to explain what value there is in coherence when it doesn't result from evidence-sensitive dispositions. Otherwise, there is no need to bifurcate structural and substantive requirements in order to explain the intuitive normative difference between subjects who differ in coherence while failing to respect their evidence. Of course, this is not the only argument for bifurcationism. In the next two sections, I'll address the argument that we should divorce substantive and structural requirements because they can come into conflict when you have evidence that supports incoherent beliefs.

4. The Unificationist Solution

According to my unificationist proposal, there is just one fundamental requirement of epistemic rationality – namely, the evidentialist requirement to proportion your beliefs to your evidence. Your beliefs are proportioned to your evidence when your degree of belief in any given proposition matched the degree to which your evidence supports that proposition. Moreover, the degree to which your evidence supports a proposition is a function of two things: substantive facts about your evidence together with structural facts about the evidential support relation. Hence, the evidentialist requirement incorporates both structural and substantive dimensions: epistemic rationality requires that your beliefs cohere with substantive facts about your evidence in accordance with structural facts about the evidential support relation.

How does this version of unificationism solve our puzzle about the conflict between evidence and coherence? On this view, there can be no conflict between evidence and coherence (cf. Kolodny 2007; Kiesewetter 2017; Lord 2018). Anyone who respects their evidence is guaranteed to be coherent, since there are coherence constraints built into the structure of the evidential support relation. These structural constraints on the evidential support relation guarantee that your evidence never supports incoherent beliefs. We can therefore maintain that

epistemic rationality always requires not only coherence, but also respecting your evidence, without countenancing epistemic dilemmas.

Formal theories of epistemic rationality typically impose *logical constraints* on the evidential support relation. These constraints explain the plausible datum that epistemically rational agents are logically coherent. Epistemically rational agents have logically consistent beliefs because they always believe what their evidence supports and their evidence never supports logically inconsistent propositions. Similarly, they believe all the logical consequences of their beliefs because the evidential support relation is closed under logical consequence:

The Evidential Closure Principle: Necessarily, if p entails that q , and your total evidence supports p to degree n , then your total evidence supports q to degree n or greater.

This principle implies that every logical truth is supported to the maximal degree – namely, certainty – by every possible body of evidence. After all, any logical truth is entailed by anything else and entailment is the strongest kind of support relation. Hence, this principle encodes a requirement of *logical omniscience*: since all logical truths are certain on your evidence, epistemic rationality requires that you should be certain of any logical truth towards which you adopt any doxastic attitude at all.

Some epistemologists take lottery and preface paradoxes to undermine logical consistency and closure requirements on epistemic rationality. However, the logical constraints on belief can be preserved in the form of probabilistic constraints on degrees of belief or credences (Christensen 2004). On this view, epistemic rationality requires that your credences are probabilistically coherent in the sense that they conform to the axioms of the probability calculus:

- (1) For every p , $\Pr(p) \geq 0$.
 - (2) If p is a tautology, then $\Pr(p) = 1$.
 - (3) If p and q are mutually exclusive, then $\Pr(p \vee q) = \Pr(p) + \Pr(q)$.
- (Christensen 2004: 16)

On a probabilistic conception of the evidential support relation, degrees of evidential support are evidential probabilities. Epistemically rational thinkers are probabilistically coherent because the evidential support relation is constrained by the axioms of the probability calculus. In particular, epistemically rational thinkers are logically omniscient because it is an axiom that logical truths always have probability 1. Hence, probabilistic coherence also encodes a requirement of logical omniscience.

In addition to these logical or probabilistic constraints, we should recognize higher-order constraints on the evidential support relation. On a probabilistic conception of evidential support, these can be formulated as constraints on higher-order probabilities, such as the following:

Probabilistic Accessibilism: Necessarily, if it is evidentially probable that p to degree n , then it is evidentially certain that it is evidentially probable that p to degree n (Smithies 2019: 230).

The rationale for higher-order constraints is to explain why epistemic rationality requires meta-coherence. Intuitively, epistemically rational agents always believe what they believe they should believe. Just as it seems irrational to act akratically in conflict with your beliefs about how you should act, so it seems irrational to believe akratically in conflict with your beliefs about how you should believe. To explain why epistemic akrasia is always irrational, we need to recognize higher-order constraints as well as first-order logical or probabilistic constraints on the evidential support relation. Epistemically rational agents are meta-coherent because they always proportion their beliefs to their evidence and they always know with certainty what their own evidence supports.

This higher-order constraint on the evidential support relation is extremely demanding, but I doubt we can settle for anything weaker. For example, Adam Elga's (2013) New Rational Reflection Principle doesn't require being certain of the evidential probability that p , so long as your credence in p matches the weighted average of your expectations about the evidential probability that p . And yet this principle is not strong enough to prohibit a form of epistemic akrasia in which you are certain that your credence is irrational, although you have no idea whether it should be higher or lower. To rule this out, we need to maintain that evidential probabilities are always evidentially certain.

This higher-order constraint on the evidential support relation can be explained as a consequence of two more basic assumptions. First, necessary truths about the evidential support relation have the same epistemic status as necessary truths about logic. Just as logical truths are certain given any possible body of evidence, so are necessary truths about the evidential support relation. The normalization axiom assigns probability 1 to all necessary truths that hold throughout the epistemic space over which evidential probabilities are defined. These epistemic necessities include necessary truths about the evidential support relation as well as necessary truths about logic. This yields an evidentialist version of Titelbaum's *fixed-point thesis*, according to which "no situation rationally permits an a priori false belief about which overall states are rationally permitted in which situations" (2015: 293).

Second, all contingent truths about your evidence are *self-evident* in the sense that they make themselves evident:

The Self-Evidence of Evidence: Necessarily, if your evidence includes (or excludes) the fact that p , then it's evidentially certain that your evidence includes (or excludes) the fact that p .

The claim that all evidence is self-evident is a plausible consequence of a phenomenal conception of evidence, according to which your evidence is

exhausted by phenomenally individuated facts about your current mental states. On this view, you have the same evidence as your phenomenal duplicate who is deceived by an evil demon, since there is no difference in how things seem to you. Assuming that skepticism is false, your evidence that it seems to you that p favors the anti-skeptical hypothesis that p over the skeptical hypothesis that it merely falsely seems that p . Arguably, however, since your evidence is consistent with the skeptical hypothesis, it doesn't rule it out with certainty, but merely with a high degree of probability.

In contrast, propositions about your own phenomenal evidence are immune from demonic deception. Your phenomenal evidence can be misleading about how things *are* but not about how things *seem*. This is because your evidence about how things seem is constituted by the facts about how things seem, rather than by second-order seemings that can misrepresent those phenomenal facts. When it seems that p , it is evidentially certain that it seems that p , since your evidence is inconsistent with any skeptical possibility in which things seem otherwise. A demon can induce false beliefs about how things seem, but he cannot induce justified false beliefs by giving you misleading evidence. Your evidence never justifies false beliefs about how things seem, since your evidence about how things seem is constituted by how things seem. Your phenomenal evidence is self-evident in the sense that it entails itself and thereby makes itself certain.

With these two claims in hand, we can explain why your evidence always makes it certain whether it supports any given proposition to any given degree. This is because contingent truths about your evidence and necessary truths about the evidential support relation are always certain given your evidence. Necessarily, if your total evidence e makes it evidentially probable for you that p to degree n , then it is evidentially certain for you that:

- (1) You have total evidence e .
- (2) If you have total evidence e , then it is evidentially probable for you that p to degree n .
- (3) Therefore, it is evidentially probable for you that p to degree n .

In sum, higher-order constraints on the evidential support relation can be explained as a natural consequence of a phenomenal conception of evidence, which can be motivated on independent grounds by appealing to standard internalism intuitions about skeptical scenarios.

It is a familiar claim that epistemic rationality requires *logical omniscience*, but this can be regarded as a special case of the more general claim that epistemic rationality requires *evidential omniscience*. Perfectly rational agents are not only certain of all logical truths but they are also certain of all truths about what their evidence is and what it supports. This requirement may seem unduly demanding, but it is a consequence of the plausible thesis that epistemic rationality requires respecting logic and evidence while remaining meta-coherent. If you violate the requirement of logical or evidential omniscience, and you integrate your reasoning

with your beliefs about logic and evidence, then your reasoning fails to respect logic and evidence. Your doubts about logic and evidence “trickle down” in ways that lead you to disrespect logic and evidence. These requirements are not mere scientific idealizations – that is, false predictions of a theory that can be safely ignored for practical purposes. Rather, they are normative ideals that non-ideal agents can approximate towards, although we can never realize them perfectly.

The view outlined in this section is developed in greater depth and detail elsewhere (Smithies 2019). My main aim here is just to explain how this view solves our puzzle by precluding conflicts between evidence and coherence. The challenge that remains is to explain away the intuitions about cases that generate the puzzle in the first place. I’ll address this challenge by invoking a distinction between ideal and non-ideal requirements of epistemic rationality.

5. Ideal and Non-Ideal Rationality

As we saw in §1, Worsnip’s (2018a) argument for bifurcationism assumes that you can have misleading higher-order evidence about what your evidence supports. For example, Miss Marple’s expert testimony gives Mabel misleading higher-order evidence that her evidence incriminates the vicar. I deny this assumption. You cannot have misleading higher-order evidence about what your evidence supports because these facts are always certain given your evidence.

We need to rethink the assumption that you have evidence for a conclusion whenever someone credible tells you that it’s true. What your evidence supports is a matter that depends not only on substantive facts about what evidence you have, but also on structural facts about the evidential support relation that apply to everyone. The structural constraints on the evidential support relation guarantee that your evidence cannot be misleading about logic and evidence. All contingent truths about your evidence are self-evident in the sense that they make themselves certain. Moreover, all necessary truths about logic and evidential support are certain given any possible body of evidence. This means, ironically enough, that when you receive misleading testimony that not- p , your evidence makes it certain that p , when this is a truth about logic or evidence.

Much of the resistance to this proposal stems from the following argument, although it is often left implicit:

- (1) Epistemic rationality always requires that you proportion your beliefs to your evidence.
- (2) When you receive misleading testimony, epistemic rationality sometimes requires that you are uncertain or mistaken about logic and evidence.
- (3) Therefore, misleading testimony sometimes provides you with uncertain or misleading evidence about logic and evidence.

As I’ll explain, however, this argument trades on an equivocation between ideal and non-ideal standards of epistemic rationality. By ideal standards, epistemic rationality always requires proportioning your beliefs to your evidence. Since your

evidence is never misleading about logic and evidence, however, it is never ideally rational to be uncertain about logic or evidence. On this reading, premise (1) is true but premise (2) is false. By non-ideal standards, in contrast, epistemic rationality sometimes requires being uncertain or mistaken about logic and evidence. But this is because non-ideal rationality requires responding to evidence about your cognitive limitations by adopting strategies that diverge from the epistemic ideal of respecting your evidence. On this reading, premise (2) is true but premise (1) is false. The argument is unsound because there is no consistent interpretation on which both premises are true.

Here is a simple example to illustrate the distinction between ideal and non-ideal standards of epistemic rationality. Suppose Holmes and Watson are working on a murder case and share all their evidence. The evidence is complicated enough that it's not obvious which conclusion it supports, but in fact it incriminates the butler. Since Holmes is an expert detective, he knows that the butler is guilty, although he keeps his opinion private. Meanwhile, Watson doesn't know what to make of the evidence, since this is not his field of expertise.

What should Watson believe? There is no single answer to this question. Deontic terms are highly context-sensitive and we get different answers depending on how we interpret them. One dimension of this context-sensitivity concerns the distinction between more and less demanding normative standards.

In one sense, Watson should believe what Holmes believes, since this is the conclusion that is supported by his evidence. Although Holmes and Watson differ in their response to their evidence, there is no difference in what their evidence supports. After all, they have exactly the same evidence. Moreover, there can be no difference in what this shared body of evidence supports, since the evidential support relation applies to everyone in the same impersonal way. Hence, Watson should believe what his evidence supports – namely, that the butler is guilty.

In another sense, however, Watson should remain agnostic. After all, he knows that – unlike Holmes – he lacks the expertise to follow the evidence where it leads. If he gets lucky in this case, he is prone to go awry elsewhere, since he is not reliably responsive to the facts or the evidence in hard cases. So, even if he forms a true belief that is supported by his evidence, he is not reliable enough to acquire knowledge or justified belief. Moreover, Watson has enough evidence about his own cognitive limitations to know all this. Since he knows he cannot acquire knowledge or justified belief, it doesn't make sense for him to try. Instead, it makes sense to adopt a more cautious policy that takes his evidence about his cognitive limitations into account. Hence, Watson should remain agnostic, rather than forming any opinion about the case.

These two answers reflect the distinction between ideal and non-ideal standards of epistemic rationality. Ideal standards of epistemic rationality always require respecting your evidence, whereas non-ideal standards sometimes require responding to evidence about your cognitive limitations by adopting strategies that diverge from the epistemic ideal. Our intuitive judgments about what people

“should” believe don’t always track what their evidence supports, since they are often more sensitive to non-ideal standards of epistemic rationality.

Someone might protest that Holmes and Watson don’t share the same evidence, since they have different evidence about their own expertise. This is true, of course, but it doesn’t mean they have different evidence about the murder case. The higher-order evidence doesn’t change the evidential probability that the butler is guilty. The evidential probability of a hypothesis depends on how well the hypothesis explains all the evidence. This is an objective, a priori matter that is not affected by evidence about your own capacity for reasoning. The evidential probability of a hypothesis isn’t affected by the realization that you’re too tired or distracted to reason clearly. This distorts the epistemic function of higher-order evidence about your cognitive limitations (Christensen 2010: 203-4).

I propose an alternative account of the epistemic function of higher-order evidence (cf. Smithies 2019: Ch. 10; forthcoming). Rather than changing what your evidence supports, it changes which response to your evidence is required by non-ideal standards of epistemic rationality. When you have higher-order evidence about your cognitive limitations, it makes sense to adopt epistemic policies to manage them. For example, it makes sense for Watson to remain agnostic, rather than forming an opinion about the case, since he knows he cannot reliably follow the evidence where it leads. It doesn’t follow that his evidence supports agnosticism. This is simply the best epistemic policy for managing what he knows about his on cognitive limitations.

We can make this proposal more precise by locating it within the framework of *rule consequentialism*, which evaluates rules by their expected consequences. We can evaluate rules in a way that is sensitive to the distinction between *following* a rule and merely *trying* to follow a rule. Following a rule is a kind of achievement: merely trying to follow the rule does not guarantee that you will succeed. When you have evidence that you might fail, the expected consequences of trying to follow a rule can diverge from the expected consequences of following the rule. In such cases, the best rule to follow is not always the best rule to try to follow (cf. Lasonen-Aarnio 2010; Schoenfield 2015).

When we evaluate rules for epistemic rationality, we’re concerned solely with their expected consequences for how well you succeed in proportioning your beliefs to your evidence. From an evidentialist perspective, the best rule to follow is *the evidentialist rule*, “Always proportion your beliefs to your evidence!” However, this is not always the best rule to try to follow when you have evidence that you are likely to fail. It is counterproductive to try to follow the evidentialist rule when this is likely to make you less responsive to your evidence. In such cases, there may be greater expected value in trying to follow some alternative strategy. By ideal standards, epistemic rationality always requires following the evidentialist rule, since this is the best rule to follow. By non-ideal standards, however, epistemic rationality sometimes requires following a non-evidentialist rule when this is the best rule to try to follow.

Now let's apply this distinction to our example. By ideal standards of epistemic rationality, Watson should believe what his evidence supports – namely, that the butler is guilty and that his evidence supports this conclusion. And yet Watson is a non-ideal agent who is always capable of achieving these demanding epistemic standards. In hard cases, he is unable to follow his evidence where it leads. Moreover, he has enough higher-order evidence about his own cognitive limitations to know this about himself. So he knows that it's counterproductive to try to respect his evidence in hard cases, since the expected consequence is that he will manifest grossly irrational dispositions. It makes more sense to adopt the cautious epistemic policy of remaining agnostic in hard cases, although he knows in advance that this strategy will diverge from the epistemic ideal. Nevertheless, adopting this strategy is a reasonable response to his higher-order evidence about his cognitive limitations.

The key point is that our intuitions about what people “should” believe don't always track what their evidence supports. There is an intuitive sense in which Watson should remain agnostic about the first-order question of whether the butler is guilty and the higher-order question of what his evidence supports. However, it doesn't follow that his evidence supports agnosticism about either first-order or higher-order questions. It's easy to overlook this point unless we pay careful attention to the distinction between ideal and non-ideal standards of epistemic rationality.

6. Epistemic Dilemmas?

Are there epistemic dilemmas in which the requirements of ideal and non-ideal rationality come into conflict? The whole point of distinguishing these requirements is that they can diverge, since ideal rationality always requires respecting your evidence, whereas non-ideal rationality sometimes requires disrespecting your evidence. For example, Watson is required by ideal standards to believe what his evidence supports – namely, that the butler is guilty – although he is required by non-ideal standards to remain agnostic. This is not an epistemic dilemma in the strict sense, however, since there is no univocal sense in which Watson ought and ought not to believe this conclusion. There's one sense in which he ought to believe it and another sense in which he ought to withhold belief. There are no epistemic dilemmas in which you ought in the same sense to pursue logically incompatible options.

This bears comparison with Worsnip's (2018a) *equivocation strategy* for solving our puzzle. He avoids epistemic dilemmas by denying that there is any single sense in which you ought to respect your evidence and to remain coherent. Unlike Worsnip, however, I reject the bifurcationist assumption that your evidence can support incoherent beliefs. On my unificationist proposal, there is no conflict between evidence and coherence, or between substantive and structural requirements, but only between ideal and non-ideal requirements. Ideal rationality always requires both coherence and respecting your evidence, whereas non-ideal

rationality sometimes requires violating these ideals. Let me close with some reasons for preferring this view.

First, unificationism preserves the attractive idea that there is a unified virtue of epistemic rationality that incorporates both substantive and structural dimensions. As we noted at the outset, epistemically rational thinkers are not only coherent but also respect their evidence. On the unificationist view, this is no mere coincidence, since respecting your evidence guarantees coherence. The virtue of epistemic rationality is to hold beliefs that cohere with substantive facts about your evidence in accordance with structural facts about the evidential support relation. According to bifurcationism, in contrast, there are two distinct virtues corresponding to the distinction between substantive and structural requirements, but there is nothing that unifies them in a single virtue of epistemic rationality. We might decide to call someone 'epistemically rational' only if they satisfy both kinds of requirements. But this doesn't pick out any unified virtue, as opposed to a gerrymandered conjunction of distinct virtues. Decomposing epistemic rationality in this way seems like a theoretical last resort.

Second, we do not compromise the unity of epistemic rationality by drawing a distinction between ideal and non-ideal requirements. On the version of rule consequentialism outlined in §5, non-ideal rationality is a matter of adopting strategies that have the greatest expected value when evaluated by standards of ideal rationality. Hence, non-ideal rationality is explained in terms of its conduciveness towards ideal rationality. Ideal and non-ideal requirements of epistemic rationality ultimately flow from the same normative source. According to bifurcationism, in contrast, substantive and structural requirements are two distinct and *sui generis* sources of normativity.

Third, unificationism is more parsimonious. Every normative theory needs some version of the ideal/non-ideal distinction to account for cases in which the expected value of trying to follow its requirements diverges from the expected value of successfully following them. In such cases, there is an ideal sense in which you should follow its requirements, but there is a non-ideal sense in which you should do otherwise when this has greater expected value. Bifurcationism doesn't obviate the need for this distinction: if we divorce substantive and structural requirements, then we need to distinguish further between ideal and non-ideal species of each genus. Hence, unificationism retains the advantage of theoretical parsimony. We cannot simply trade the distinction between ideal and non-ideal requirements for the distinction between substantive and structural requirements. These distinctions don't do the same kind of theoretical work.

Fourth, unificationism explains how there can be epistemic value in disrespecting your evidence. As we've seen, bifurcationism struggles to explain the value of coherence when it comes at the cost of respecting your evidence. In contrast, unificationism does a better job of explaining the value of non-ideal rationality when it diverges from the epistemic ideal of respecting your evidence. There is value in non-ideal rationality because it maximizes your expected degree of responsiveness to your evidence given your higher-order evidence about your

cognitive limitations. Hence, the value of non-ideal rationality can be explained in evidentialist terms.

Finally, unificationism explains our intuitions about cases with minimal mutilation of plausible theoretical principles. We can explain the intuitive sense in which Watson should remain agnostic about the murder, despite the fact that his evidence incriminates the butler. This intuition tracks a non-ideal requirement of epistemic rationality, which permits Watson to respond to higher-order evidence about his own cognitive limitations by adopting strategies that deviate from the epistemic ideal. At the same time, however, we can retain the plausible theoretical principle that epistemic rationality always requires respecting your evidence, since evidentialism is true of the ideal requirements of epistemic rationality. In this way, the distinction between ideal and non-ideal rationality allows us to reconcile intuition and theory without making recourse to bifurcationism.

References

- Christensen, David. 2004. *Putting Logic In Its Place: Formal Constraints on Rational Belief*. Oxford University Press.
- Christensen, David. 2007. Does Murphy's Law Apply in Epistemology? Self-Doubt and Rational Ideals. *Oxford Studies in Epistemology* 2: 3-31.
- Christensen, David. 2010. Higher-Order Evidence. *Philosophy and Phenomenological Research* 81.1: 185-215.
- Coates, Allen. 2012. "Rational Epistemic Akrasia." *American Philosophical Quarterly* 49.2: 113-124.
- Elga, Adam. 2013. "The Puzzle of the Unmarked Clock and the New Rational Reflection Principle." *Philosophical Studies* 164.1: 127-139.
- Enoch, David. 2006. Agency, Schmagency: Why Normativity Won't Come From What is Constitutive of Action. *Philosophical Review* 115.2: 169-198.
- Feldman, Richard. 2003. *Epistemology*. Prentice Hall.
- Horowitz, Sophie. 2014. Epistemic Akrasia. *Nous* 48.4: 718-44.
- Huemer, Michael. 2001. *Skepticism and the Veil of Perception*. Rowman and Littlefield.
- Hughes, Nick. 2019. Dilemmic Epistemology. *Synthese* 196.10: 4059-4090.
- Kiesewetter, Benjamin. 2017. *The Normativity of Rationality*. Oxford University Press.
- Kolodny, Niko. 2005. Why Be Rational? *Mind* 114: 509-563.
- Kolodny, Niko. 2007. How Does Coherence Matter? *Proceedings of the Aristotelian Society* 107: 229-263.
- Lasonen-Aarnio, Maria. 2010. Unreasonable Knowledge. *Philosophical Perspectives* 24: 1-21
- Lasonen-Aarnio, Maria. 2020. Enkrasia or Evidentialism? Learning to Love Mismatch. *Philosophical Studies* 177: 597-632.
- Leonard, Nick. 2020. Epistemic Dilemmas and Rational Indeterminacy. *Philosophical Studies* 177: 573-596.
- Lord, Errol. 2018. *The Importance of Being Rational*. Oxford University Press.

- Sartre, Jean-Paul. 1946. *Existentialism is a Humanism*. Methuen.
- Scanlon, T. M. 1998. *What We Owe To Each Other*. Harvard University Press.
- Scanlon, T. M. 2007. Structural Irrationality. In *Common Minds: Themes From the Philosophy of Philip Pettit*, edited by Geoffrey Brennan, Robert Goodin, Frank Jackson, and Michael Smith. Oxford University Press.
- Schoenfield, Miriam. 2015. Bridging Rationality and Accuracy. *Journal of Philosophy* 112.12: 633-657.
- Schroeder, Mark. 2008. Having Reasons. *Philosophical Studies*. 139: 57-71.
- Smithies, Declan. 2012. Moore's Paradox and the Accessibility of Justification. *Philosophy and Phenomenological Research* 85.2: 273-300.
- Smithies, Declan. 2019. *The Epistemic Role of Consciousness*. Oxford University Press.
- Smithies, Declan, Preston Lennon, and Richard Samuels. Forthcoming. Irrational Minds. *Synthese*.
- Titelbaum, Michael. 2015. Rationality's Fixed Point (Or: In Defense of Right Reason). *Oxford Studies in Epistemology* 5: 253-294.
- Weatherson, Brian. 2019. *Normative Externalism*. Oxford University Press.
- Wedgwood, Ralph. 2017. *The Value of Rationality*. Oxford University Press.
- Worsnip, Alex. 2018a. The Conflict of Evidence and Coherence. *Philosophy and Phenomenological Research* 96.1: 3-44.
- Worsnip, Alex. 2018b. What Is (In)Coherence? *Oxford Studies in Metaethics* 13: 184-206.