Imperfection, Accuracy, and Structural Rationality

Marc-Kevin Daoust, École de technologie supérieure. marc-kevin.daoust@etsmtl.ca

Paper forthcoming in *Erkenntnis*. DOI: (10.1007/s10670-022-00571-5). Please cite the published version of the manuscript.

Abstract: Structural requirements of rationality prohibit various things, like having inconsistent combinations of attitudes, having means-end incoherent combinations of attitudes, and so on. But what is the distinctive feature of structural requirements of rationality? And do we fall under an obligation to be structurally rational? These issues have been at the heart of significant debates over the past fifteen years. Some philosophers have recently argued that we can unify the structural requirements of rationality by analyzing what is *constitutive* of our attitudes (Lee 2020; Worsnip 2018a). It has also been suggested that, in the course of good first-personal deliberation, agents should treat structurally irrational combinations of attitudes as off-limits (Worsnip 2021). In this paper, I raise a worry for these two theses concerning structural rationality. Roughly, I argue that some imperfect epistemic agents (like us) can be disposed to have structurally irrational combinations of attitudes. Also, when these imperfect agents deliberate, they should not treat structurally irrational combinations of attitudes as off-limits. Given our imperfections, being structurally irrational can very well be the best option we have. More generally, these observations reveal that structural rationality should not always be theorized independently from more "substantive" norms, like responsiveness to reasons or expected value optimization.

Keywords: structural rationality, coherence, consistency, epistemic ideals, norms of perfection, bounded rationality

Suppose Sam has decisive evidence to believe that Megan is guilty. However, he disbelieves that Megan is guilty. There is a sense in which we can say that Sam is irrational. That is, he fails to respond correctly to the evidence he has. And substantive rationality requires of agents that they respond correctly to the epistemic reasons (or, in this case, the evidence) they have. So, we can say that Sam is irrational in the *substantive* sense.[1]

Now, suppose that Sam believes that Megan is guilty, but also that she is innocent. Here, Sam is also irrational, but in a different way. Not only does he fail to respond correctly to his evidence, but his attitudes are in tension with each other. The difference from the previous case is that here he also violates some putative *structural* requirements of rationality. Let's say,

1 See, e.g., Kiesewetter (2017) and Lord (2018).

provisionally, that structural requirements of rationality prohibit inconsistent combinations of beliefs or intentions, or akratic combinations of attitudes, and the like.[2] Here is a partial list of putative structural requirements of rationality discussed in the literature:

Belief Consistency. Rationality requires that, if A believes that P, A does not believe ~P.[3]

Probabilistic Consistency. Rationality requires that, if A assigns a credence of X in P, A assigns a credence of (1-X) in P.[4]

Practical Consistency. Rationality requires that, if A intends to φ, A does not intend not to φ.[5]

Instrumental Principle. Rationality requires that, if A intends to φ, and A believes that ψ-ing is a necessary means to φ-ing, then A intends to ψ.[6]

Structural requirements of rationality have been at the heart of two important debates in the past decades. First, what is the distinctive feature of structural requirements of rationality? We have a long list of putative structural requirements.[7] And of course, we can analyze each putative requirement individually, and try to figure out whether it is good, bad, required, permitted, and so on to satisfy some of these specific requirements. But a more *systematic* way to analyze structural rationality would be to work from a distinctive feature underlying all the putative structural requirements. So, can we find a common denominator underlying the putative structural requirements? Call this the Descriptive Question.[8]

Second, do we fall under an obligation to satisfy the structural requirements of rationality? At least, do we have a reason to be structurally rational? Call this the Normative Question. Many philosophers such as Kolodny (2005; 2007; 2008a; 2008b) have argued that structural requirements of rationality are not normative (i.e., that agents do not fall under an obligation to be structurally rational). Other philosophers, such as Broome (2013), believe that these requirements are normative, but can't find an explanation why.

---

2 See, e.g., Broome (2013), Kolodny (2005; 2007; 2008a; 2008b), and Worsnip (2018a; 2018b; 2019; 2021). Putative structural requirements can be found in the literature on the normativity of rationality, but also in decision theory and in formal epistemology. See, e.g., Pettigrew (2016).

3 See Broome (2013, 154–56) and Kolodny (2007).

4 See Joyce (1998; 2009) and Pettigrew (2016). Probabilistic Consistency could also be a consequence of Broome's "Bayesian Requirement"—see Broome (2013, 175).

5 See Broome (2013, 156) and Kolodny (2008).

6 See Brunero (2009; 2020), Broome (2013, sec. 9.4), Kiesewetter (2017), Lord (2014; 2018), and Kolodny (2005) on the Instrumental Principle.

7 See Daoust (2022).

8 Broome (2013, 150), for instance, analyzes putative requirements individually and is uncertain whether there is a common denominator underlying them. For him, each requirement can be justified on different grounds.

In some recent contributions, Worsnip (2018a; 2021) tries to answer both questions. In response to the first question, Worsnip offers a dispositionalist account of coherence tailored to unify the structural requirements of rationality. A similar account of coherence can also be found in Lee (2020). This provides an answer to the Descriptive Question. In response to the second question, Worsnip argues that structural requirements set the limits of first-personal deliberation. When we deliberate, we should ignore combinations of attitudes that are structurally rational or treat them as off-limits. In that regard, we should be structurally rational. This provides an answer to the Normative Question.

In this paper, I present a worry for these answers to both questions. The gist of the argument is this: Some agents (like us) are imperfect. They fail to satisfy some substantive requirements. For instance, some of their attitudes are unreasonable, or are not supported by the evidence or fail to maximize expected value. However, even if these agents fail to satisfy some substantive norms, they can still be perfectly coherent, or structurally rational. Given their imperfections, should they treat structurally irrational combinations of attitudes as off-limits in first-personal deliberation? I argue that this would be a bad policy. Imperfect agents should sometimes treat structurally irrational combinations of attitudes as legitimate options.

In section 1, I present Worsnip's answers to the Descriptive and the Normative questions. In section 2, I present the case of an Imperfect Optimizer. In section 3, I argue that the possibility of an Imperfect Optimizer causes trouble for the dispositionalist account of coherence. In section 4, I argue that, when imperfect optimizers deliberate, they should not ignore structurally irrational combinations of attitudes. I also respond to two objections against my argument. In section 5, I discuss one of the main upshots of the paper: We should not always analyze structural and substantive rationality separately from each other. If you can't be fully substantively rational, then your second-best option might be to violate the structural requirements of rationality. We should take these interactions between structural and substantive norms into account.

1. Worsnip's Answers to the Descriptive and Normative Questions

1.1. What Is Structural Rationality (or Coherence)?

Worsnip (2018a; 2018b; 2019; 2021) and Lee (2020) think that structural rationality is coherence. That is, the property underlying putative structural requirements is coherence. A similar claim can be found in Kolodny (2005, 511). But what does "coherence" mean, exactly? That is, how do we substantiate the meaning of "coherence," so that we can offer an informative answer to the Descriptive Question? This is not a trivial question. There are various theories and accounts of coherence in print: coherent agents are not vulnerable to exploitation or do not have dominated combinations of attitudes,[9] coherent agents have attitudes that mutually support each other,[10] coherent agents have attitudes that are not in tension with their phenomenal experiences,[11] coherent agents are not logically inconsistent,[12] and so on. So, structuralists about rationality should tell us what they mean by "coherence."

This is exactly what Worsnip does. He takes coherence to be a regimented concept tailored to unify structural requirements of rationality (Worsnip 2018a, 187). His regimented account of coherence states the following:

Coherence. "A set of attitudinal mental states is jointly incoherent iff it is (partially) constitutive of the mental states in question that, for any agent that holds these attitudes, the agent is disposed, when conditions of full transparency are met, to give up at least one of the attitudes." (Ibid., 188)

There are similar accounts of coherence and structural rationality in print, like Lee's (2020).[13] For simplicity and concreteness, I here focus on Worsnip's specific formulation.

There are three important concepts in Worsnip's account of coherence—namely, transparency, constitution, and dispositions. First, here is how Worsnip defines transparency:

9 See, e.g., Davidson, McKinsey and Suppes (1955), de Finetti (1937; 1974), Easwaran and Fitelson (2015), Hájek (2009), and Joyce (1998).

10 See, e.g., C. I. Lewis (1946, 338), Ekstrom (1993, 608), Raz (1992, 293), Sayre-McCord (1985, 171), and Thagard (2002, 43).

11 See Wedgwood (2017, 11-12).

12 Davidson (1985, 346) makes a similar claim.

13 Lee says that "it is partly constitutive of belief and intention that your occurrent (or activated) beliefs and intentions satisfy the coherence requirements" (Lee 2020, 6). Lee's and Worsnip's accounts of coherence both focus on what is *constitutive* of attitudes. Also, they both focus on attitudes that are somehow *present* to the agent's mind (since they are concerned with attitudes that are occurrent and/or transparent). However, Lee's account of coherence is stronger than Worsnip's. Under some conditions C, Worsnip's account says that agents are *disposed* to be coherent, while Lee's account says that agents *are* coherent.

> By "conditions of full transparency," I mean conditions under which the agent knows, and explicitly and consciously believes, that she has the states in question, without self-deception, mental fragmentation, or any failure of self- knowledge (pertaining to those attitudes). (Ibid.)

So, very roughly, an agent's attitudes are transparent when they are *known* and *conspicuous* to him or her. Then, there is constitution. For Worsnip, a set of attitudes is incoherent if it is *constitutive* that, when some transparency conditions are satisfied, agents are disposed to give at least one of them up. Worsnip doesn't define what he means by "constitutive." Plausibly, he here means something like "essential" (as opposed to "derived from what makes us agents"). Finally, there are dispositions. For Worsnip, when agents are disposed to give up some combinations of attitudes, they would find it hard to maintain all their attitudes simultaneously (ibid.).

Worsnip thinks that he can unify putative structural requirements of rationality with this account of coherence. Suppose that one's attitudes are known and conspicuous to one. Then, when one violates requirements like Belief Consistency, Probabilistic Consistency, Practical Consistency or the Instrumental Principle, one is *essentially* disposed to give up some of one's mental states. That is, one would find it very difficult to maintain such (combinations of) mental states.

Worsnip does not claim that *all* putative structural requirements have this property. But if a putative structural requirement conflicts with his account of coherence, then this might be a good reason to think that this is a "false" structural requirement (ibid., 187-8). For instance, suppose one is not essentially disposed to satisfy Modus Ponens Rationality.[14] That is, suppose one is perfectly fine with having combinations of beliefs that are not closed under entailment.[15] Then, Worsnip might argue that violating Modus Ponens Rationality can be coherent, and that it is not a genuine structural requirement.

14 Modus Ponens Rationality says this: rationality requires that, if A believes that p, and A believes that $p \rightarrow q$, then A believes that q. See Kiesewetter (2017, chap. 9) for discussion.

15 We can imagine a "Lockean" agent who thinks that violating Modus Ponens Rationality is sometimes worth it. Lockeans, for instance, think that rational agents can sometimes violate Modus Ponens Rationality. See Easwaran (2016).

1.2. Worsnip's Answer to the "Making Space" Challenge

Do we fall under an obligation to be structurally rational? According to Kolodny (2005), what matters in correct first-personal deliberation is whether our reasons support some attitudes. Why should I believe that P? It is because my reasons decisively support believing P. Why am I permitted to intend to φ? It is because I have decisive (or sufficient) reasons to φ. In first-personal deliberation, it seems that all that matters is whether our reasons support our attitudes. For Kolodny, considerations of structural rationality (or coherence) do not show up in correct first-personal deliberation about what we ought to do or believe.

If we take this observation seriously, we should doubt that, in addition to being reasonable, we fall under a distinct obligation to be structurally rational. So, for those who think that structural rationality is normative, one challenge is to accommodate this datum concerning correct first-personal deliberation. This is what Worsnip (2021) calls the challenge of "making space" for the normativity of structural rationality.

Worsnip's response to this challenge goes as follows. First, he challenges the assumption that, if we fall under an obligation to be coherent, then considerations of coherence will show up in favour of particular attitudes (ibid., 7). Worsnip does not think that this is the only way coherence can be relevant and normative in first-personal deliberation. That is, he wants to "explore whether there's a different way of fitting coherence considerations into deliberation" (ibid.).

Worsnip notes that deliberation is often holistic. That is, we often plan and deliberate on various options simultaneously. As he says, "when we deliberate, we often don't deliberate about individual attitudes in isolation" (ibid., 9). For instance, my plan to drink this evening is not independent of my plan to drive after the party. My commitment to driving my car after the party matters for whether I should drink this evening. And so, when I deliberate on how I will spend my evening, I need to consider my various options together, not in isolation from each other.

In light of this, Worsnip argues that structural requirements govern our deliberation in a certain way. They rule out some combinations of attitudes (or options) as off-limits. Just as I should not commit to drinking and driving simultaneously tonight, I should not decide to intend

to φ and intend not to φ simultaneously. This combination of attitudes should be precluded in first-personal deliberation (ibid., 11). He says:

> [W]e should treat incoherent combinations of attitudes as off-limits in our deliberation, and focus our deliberative attention on the merits of the remaining, coherent combinations. If one were to deliberate in a way that took some incoherent combination of attitudes seriously as an option, then one's deliberation would be, to that extent, faulty. (Ibid.)

So, coherence is normative in the following sense: *prima facie*, you should not entertain the possibility of having structurally irrational combinations of attitudes. This provides an answer to the objection raised by Kolodny. Considerations of coherence might not show up in first-personal deliberation in favour of some specific attitudes. But they can, and should, show up for ruling out some combinations of attitudes.

2. The Imperfect Optimizer

In the next three sections, I raise a worry for the above answers to the Descriptive and the Normative questions. For concreteness, I begin by presenting cases of "Imperfect Optimizers" (§2). Then, I argue that the possibility of an Imperfect Optimizer causes trouble for the dispositionalist account of coherence (§3). I also argue that Imperfect Optimizers should not exclude incoherent combinations of attitudes from their deliberation (§4).

2.1. A Very General Description of Imperfect Optimizers

Many epistemic norms discussed in print are, to some extent, idealized. For instance, we often hear that agents fall under an epistemic obligation to satisfy the axioms of probability,[16] to make probabilistic inferences in accordance with Conditionalization,[17] to assign a credence of 1 to *a priori* truths,[18] to be immodest,[19] and to satisfy various substantive requirements governing credences, like the Principal Principle or the Principle of Indifference.[20] Surely, perfect epistemic cherubs can satisfy such requirements. But for agents like us, satisfying all these requirements

16 Joyce (1998), Leitgeb and Pettigrew (2010a; 2010b), Pettigrew (2016).
17 Greaves and Wallace (2006), Meacham (2015; 2016), Schoenfield (2017).
18 Easwaran (2011), Dogramaci (2018).
19 D. Lewis (1971), Joyce (2009), Horowitz (2014; 2019), Elga (2010; 2013), Lasonen-Aarnio (2015).
20 Pettigrew (2016, chaps. 8-10).

can be extremely demanding. We are not in a position to satisfy them (or, at least, we cannot satisfy all of them). I will focus on the case of an imperfect agent called Anna who has to deal with the fact that she can't satisfy all the demands of rationality.

There are different ways in which Anna can be less than ideal.[21] Perhaps Anna's mind is opaque or fragmented. Or perhaps Anna has limited cognitive resources for processing her evidence. For simplicity, we will assume that Anna's mind is fully transparent. Her main imperfection is that some of her attitudes violate substantive requirements (in section 2.5, we will make a slightly different assumption). We can also assume that she knows the constraints she faces.

Anna is imperfect. However, we can assume that she is disposed to do what is optimal *relative to the constraints she faces*.[22] She cannot reach the epistemic ideal, but she can still take attitudes that are "second best." I take it that most of us are like Anna. We violate *at least one* of the idealized epistemic requirements. But we still try to do the best we can in our imperfect circumstances. Call Anna an Imperfect Optimizer.

Suppose that substantive rationality gives us a ranking of psychological states, from best to worst:

Ranking of states: $S_1 > S_2 > S_3 > ... > S_n$

If we were ideal, we would take S_1. Unfortunately, this state is unavailable to us. Suppose S_2 is the substantively best state available to us. Agents like us (or like Anna) are imperfect optimizers. We are disposed to do the best we can. So, we are disposed to take S_2. However, perhaps S_2 does not satisfy the requirements of structural rationality (say that only S_1 and S_3 are structurally rational states). So, we are sometimes disposed to have structurally irrational states, and we do not always exclude them from our first-personal deliberation.

This is a very abstract description of the case. And the description leaves a lot of details open, like the following: Where does the ranking come from? Why assume that the second-best state is structurally irrational? I now want to give concrete examples illustrating this. We can

21 For instance, compare Simon (1976), Tversky and Kahneman (1983) and Staffel (2019).
22 See, e.g., Staffel (2019, esp. 112-116).

disagree with the details of some of the cases, but I am confident that readers will be satisfied with at least one of the illustrations I provide.

2.2. Probabilistic Consistency and Expected Accuracy Optimization Under Constraints

Ideal agents satisfy various epistemic norms. But where do these norms come from? According to a popular answer, substantive epistemic norms come from *expected accuracy maximization*. For instance, suppose that the Principal Principle[23] is a substantive epistemic norm. Suppose also that Anna knows that P's objective probability is 0.6. Then, the credence in P that maximizes expected accuracy (relative to the objective probability function) is 0.6.[24] So, if Anna were perfectly substantively rational, she would have a credence of 0.6 in P.

As I said earlier, Anna is imperfect. Say, her credence in P can be only equal to or greater than 0.8, and so she can't have a credence of 0.6 in P. Anna is epistemically imperfect, in the sense that her attitudes depart from the epistemic ideal. She violates a substantive requirement governing the relationship between her credences and her knowledge (or reasons).

Anna fails to maximize expected accuracy (or to minimize expected inaccuracy). Yet, she can optimize expected accuracy *under constraints*.[25] For instance, suppose Anna is trying to figure out what credences in P and in ~P she should have. She knows that P's objective probability is 0.6. But, as I said before, she unfortunately has a credence of at least 0.8 in P, and she can't change that. Still, she can optimize expected inaccuracy under constraints.

Suppose Anna optimizes her credence function under the constraint "my credence in P is equal to or greater than 0.8." What credence in ~P should she take? If she maximizes expected accuracy under constraints, she will take a credence of 0.4 in ~P and a credence of 0.8 in P.[26] Given the constraints she faces, this is the best she can do. The probabilistically consistent

---

23 The Principal Principle roughly says that an agent's credences should match his or her knowledge of the objective probabilities. For instance, if you know that P's objective probability is 0.5, then you should have a credence of 0.5 in P.

24 Relative to a scoring rule, like the Brier score. See Pettigrew (2016, ch. 8).

25 See, e.g., Staffel (2019, 112-116).

26 Relative to the Brier score, her constrained expected inaccuracy function goes as follows:

$$Constrained\ expected\ inaccuracy = 0.6\cdot(1-a)^2+0.4\cdot a^2+0.4\cdot(1-b)^2+0.6\cdot b^2, a\geq 0.8.$$

Here, a denotes Anna's credence in P, and b denotes Anna's credence in ~P. We can find the find the credence function that optimizes expected accuracy with WolframAlpha, by running the following equation: Minimize 0.6*(1-a)^2+0.4*(a^2) + 0.4*(1-b)^2 + 0.6*(b^2), a>=0.8

combinations of credences available to her, like a credence of 0.2 in ~P and a credence of 0.8 in P, would fail to maximize expected accuracy under constraints.[27]

So, Anna is disposed to maximize expected accuracy under constraints. She is imperfect because she violates a substantive requirement. Given her imperfections, Anna should *not* satisfy requirements like Probabilistic Consistency. If she limited her deliberation to probabilistically consistent combinations of credences, she would sometimes fail to optimize expected accuracy under constraints. Of course, Anna *could* choose credences that are probabilistically consistent. For instance, she *could* have a credence of 0.2 in ~P and a credence of 0.8 in P. But she has a (legitimate) preference for expected accuracy under constraints. So, she chooses to have inconsistent credences.

2.3. Probabilistic Consistency and Duties to Respect the Evidence

The previous case presupposes that substantive rationality is tied to expected accuracy optimization. Maybe you disagree with this (or with epistemic consequentalism in general!). Fortunately, we can think of a nonconsequentialist version of the argument.

Suppose, for instance, that substantively rational agents *respect* their evidence.[28] That is, agents like Anna have a duty to apportion their doxastic attitudes to the evidence available to them. Suppose that Anna's mind is transparent to her, that she cares about whether P, and that her evidence includes the fact that P's objective probability is 0.6. Then, if Anna respects her evidence, she should have a credence of 0.6 in P, and a credence of 0.4 in ~P.

Unfortunately, Anna is imperfect, and her credence in P is equal to or greater than 0.8. But she has full control over her credence in ~P. What should Anna do? Anna has two options:

1. She can *partially* respect her evidence and take a credence of 0.4 in ~P (and a credence in P located in the interval [0.8, 1]). However, Anna will then be probabilistically inconsistent.

27 Given Anna's knowledge, the expected inaccuracy of a credence of 0.2 in ~P and a credence of 0.8 in P is 0.56. The expected inaccuracy of a credence of 0.4 in ~P and a credence of 0.8 in P is 0.52.

28 See, e.g., Feldman (2005).

2. She can be probabilistically consistent. So, if she takes a credence of 0.8 in P, she takes a credence of 0.2 in ~P. If she takes a credence of 0.9 in P, she takes a credence of 0.1 in ~P, and so forth. However, her credence in ~P will not respect her evidence.

Can (and *should*) Anna be disposed to take a credence in ~P that partially respects her evidence? Recall that Anna is disposed to do what's best given the limitations she faces. So, assuming that what's best for Anna is to partially respect her evidence, she will take a credence of 0.4 in P. Once again, Anna will end up with a structurally irrational combination of attitudes.

2.4. Inter-Level Coherence and Duties to Respect the Evidence

The two previous cases presuppose that Probabilistic Coherence is a genuine structural requirement of rationality. You might disagree with this assumption. I can offer an illustration that focuses on another putative structural requirement. Let's assume, for the sake of the argument, that Inter-Level Coherence is a genuine structural requirement. So, structurally rational agents refrain from having the following combination of belief-type attitudes: suspending judgment on whether P and believing that they have decisive reasons to believe P.

Suppose once again that Anna has a duty to apportion her doxastic attitudes to the evidence available to her. Anna's evidence decisively supports believing that she has decisive reasons to believe P and believing P. Unfortunately, Anna is imperfect: P is simply *unbelievable* to her. So, she suspends judgment on whether P. Yet, Anna has control over her higher-order attitude in the proposition "I have decisive reasons to believe P." What should Anna do? Anna has two options:

1. She can *partially* respect her evidence and believe that she has decisive reasons to believe P. However, Anna will then violate Inter-Level Coherence.
2. She can satisfy Inter-Level Coherence and not form the belief that she has decisive reasons to believe P. However, her higher-order attitude will not respect her evidence.

Can (and should) Anna be disposed to believe that she has decisive reasons to believe P? It's plausible that she can and should. Since Anna is disposed to do what's best given the limitations she faces, she will partially respect her evidence. Accordingly, she will believe that

she has decisive reasons to believe P. But then, her combinations of attitudes will be structurally irrational.

2.5. An Indirect Case: Mental Processes and the Preservation of Coherence

In the above cases, Anna faces a direct constraint on the specific attitudes she can have. But the constraint Anna faces could be indirect. For instance, the constraint could target the mental processes that Anna can take.[29] Suppose that Anna has to choose a *mental process* for forming and revising attitudes. The ideal process is conducive to accuracy and guarantees coherence. Unfortunately, the ideal process is unavailable to Anna. Specifically, Anna has to make a choice between two mental processes:[30]

1. Process 1 has higher expected accuracy than Process 2, but is less coherent than Process 1 (i.e., Process 1 leads to more cases of incoherent combinations of attitudes than Process 2).

2. Process 2 has lower expected accuracy than Process 1, but is more coherent than Process 2.

For instance, suppose that Process 1 relies on semilexicographic choice and that Process 2 relies on lexicographic choice.[31] As discussed in Thorstad (2021), epistemic processes that rely on semilexicographic choice can have higher expected accuracy than the lexicographic alternative. However, processes that rely on semilexicographic choice are less coherent than the lexicographic alternative. In other words, processes that rely on lexicographic choice are a better option for preserving coherence.

Can (and should) Anna choose Process 1, even if she knows that Process 1 is less coherent? This option seems fine given her circumstances. Anna is disposed to do what's best given the limitations she faces. Taking Process 1 is second best, and so this seems to be the process she should choose. But of course, she will more often end up with combinations of attitudes that are structurally irrational.

---

29 See, e.g., Simon (1976) on procedural rationality.

30 There could be an available mental process that guarantees coherence. For instance, there could be a skeptical mental process that says "never form beliefs." Since Anna would never have beliefs, she would never have incoherent combinations of beliefs. But of course, this mental process would not be conducive of accuracy.

31 See, e.g., Fishburn (1974), Thorstad (2021) and Tversky (1969) for examples and discussion of lexicographic and semilexicographic processes.

2.6. Taking Stock

I could give more examples illustrating the general point I am trying to make. But I think that at least one of the above illustrations will convince the reader, and so I will stop here. All the cases have the same structure:

1. Some agents are imperfect.
2. Given their imperfections, these agents sometimes have to make a choice between what is second-best for them and what is structurally rational (or, in indirect cases, what will lead these agents to be structurally rational).
3. Agents can, and should, be disposed to choose what is second-best for them.

3. The Possibility of an Imperfect Optimizer and Dispositions to Be Structurally Rational

3.1. A Worry for the Dispositionalist Account of Coherence

I take the Imperfect Optimizer to be a counterexample to the dispositionalist account of coherence. My reasoning goes as follows:

P1. The following is possible: the Imperfect Optimizer violates requirements like Probabilistic Consistency or Inter-Level Coherence (or: is disposed to violate these requirements) while her mind is fully transparent to her.

P2. For the dispositionalist, when conditions of full transparency are met, if A is not essentially disposed to give up at least one of her attitudes, then A is coherent.

P3. So, violating requirements like Probabilistic Consistency, Inter-Level Coherence, and the like can be coherent (in the dispositionalist sense).

P4. However, requirements like Probabilistic Consistency, Inter-Level Coherence, and the like are structural requirements of rationality.

C. So, the dispositionalist account of coherence does not capture structural rationality.

Dispositionalists like Worsnip can't deny P2, since this is a direct consequence of their account of coherence. P3 is entailed by P1 and P2. Finally, recall that dispositionalists want to offer an account of coherence tailored to unify the structural requirements of rationality. So, they can't accept C.

In light of this, I see two escape routes for dispositionalists. First, they could deny P1. Perhaps the scenario I describe in section 2 is impossible. For instance, perhaps the Imperfect Optimizer's mind cannot be fully transparent to her. Second, they could deny P4 and argue that, in accordance with their own account of coherence, requirements like Probabilistic Consistency or Inter-Level Coherence are not structural requirements of rationality. Third, they could argue that there are some implicit conditions of rational control in their account of coherence. In the next subsections, I analyze these escape routes.

3.2. Escape Route #1: The Case I Describe Is Impossible

Dispositionalists could first deny that it is possible, for the Imperfect Optimizer, to be disposed to violate requirements like Inter-Level Coherence or Probabilistic Consistency. For instance, suppose that Anna tries to violate Probabilistic Consistency. Perhaps Anna cannot seriously claim to have a credence of 0.4 in ~P and a credence of 0.8 in P. As Worsnip says, agents like Anna would "sound like a joke."[32]

However, I find this implausible. Surely, agents can be disposed to do their best under constraints. In mundane situations, we often have a disposition to do what's best or what's right. In the cases I describe, Anna would violate Probabilistic Consistency, Inter-Level Coherence or other requirements because violating these requirements is the best available option she has. Accordingly, she can be be disposed to have structurally irrational combinations of attitudes. Thus, we should grant that the case I describe is *possible*.

Dispositionalists could reply that Anna can be disposed to have structurally irrational combinations of attitudes, but only because she violates the transparency condition. For instance, the Imperfect Optimizer is disposed to violate Probabilistic Consistency because she is unaware of her own mental states, or has a fragmented mind. However, this is also implausible. In section 2, we assumed that Anna is perfectly aware of her own mental states. And she is even in a position to *understand why* she takes these attitudes: given her imperfections and limitations, she wants to do what's best (and she will sometimes succeed to do what's best relative to her imperfections). In view of the foregoing, P1 makes sense.

32 Worsnip (2018a, 191-2).

3.3. Escape Route #2: False Structural Requirements

What about P4? In accordance with their account of coherence, dispositionalists like Worsnip could deny that the requirements I focused on are genuine structural requirements. Once again, take Probabilistic Consistency. Dispositionalists could say: Since Anna can see that violating Probabilistic Consistency is sometimes the best option she has, then violating Probabilistic Consistency is not a genuine structural requirement of rationality. In other words, the dispositionalist account of coherence helps us see that this is a false structural requirement. And we could say the same thing about Inter-Level Coherence or other putative requirements.

There is not much I can say in response to this suggestion. All I can say is that this would be a significant shift in the discussion initiated by Kolodny and others. Virtually everyone agrees that Probabilistic Consistency and Inter-Level Coherence are structural requirements of rationality. If we deny this, it is no longer clear that we are still concerned with the same phenomenon.

For instance, take Probabilistic Consistency. A long tradition of empirical studies on rationality holds that satisfying the axioms of probability is a necessary condition for rationality.[33] Some authors, like Jeffrey (2004, 102), use the expression "structural rationality" to refer to the satisfaction of Bayesian requirements and issues of "bare logic." When Zynda (1996) says that coherence is an ideal of rationality, he is primarily concerned with probabilistic coherence. Broome (2013, 175) thinks that some Bayesian requirements, like Probabilistic Consistency, are requirements of rationality. In Kolodny's work, it is also very clear that Probabilistic Consistency is a structural requirement of rationality. In the first paragraphs of "How Does Coherence Matter," he says that there are "[rational] requirements of formal coherence...to have degrees of belief that satisfy the axioms of probability" (Kolodny 2007, 229). He also says that, when he thinks of structural requirements of rationality, austere Bayesianism is one of the theories he has in mind (ibid., 230).

Suppose we deny P4. Then, when dispositionalists say that "we should be coherent" or that "we should be structurally rational," they mean something different than Kolodny, Broome, and others do. And this strikes me as a problem for making progress in this debate, because we

33 See, e.g., Arkes et al. (2016). See Tversky and Kahneman (1983) for a classic reference.

are no longer answering the exact same question. If we want to have a real dialogue with Broome, Kolodny and others, we should grant that requirements like Probabilistic Consistency are structural requirements.

Dispositionalists could reply that the putative structural requirements discussed in print must be open to criticism. It should be possible to deny some of these requirements. To be clear: I agree that it must be possible to challenge *some* of the putative requirements discussed in the past fifteen years. There are obvious difficulties with having to infer a concept of coherence from a long list of putative structural requirements. But I think we should accept some "fixed points" in this debate. For instance, violating Probabilistic Consistency is one of the clearest examples of structural irrationality. It has been at the heart of the Why-Be-Rational? debate since the beginning. So, I think that accounts of coherence that unify structural requirements of rationality should not conflict with requirements like Probabilistic Consistency.

Here is another way to deny P4. Dispositionalists could suggest that we should relativize the structural requirements to the capacities of agents. Recall that the Imperfect Optimizer is stuck with a substantively irrational attitude. And this leads her to have a structurally irrational combination of attitudes. Here, perhaps the lesson is that structural requirements should be sensitive to the capacities of agents. For instance, Probabilistic Consistency could be revised as follows: Rationality requires that, *if A has the capacity to choose and revise his or her credences*, and A assigns a credence of X in P, then A assigns a credence of (1-X) in P.

This is an interesting suggestion. However, I have little to say in response to it. Again, if this were the right solution to the puzzle, we would shift the discussion initiated by Kolodny and others. If we look at the putative structural requirements discussed in print (in, e.g., Kolodny's or Broome's work), none of them refers to the capacities of agents.[34] So, perhaps we should revise all the putative structural requirements, so that they refer to the capacities agents have. But this would be a significant turn in the debate. We would end up unifying some *reformulated* versions of the putative structural requirements. And, once again, this would be an obstacle to making progress in this debate, because we would no longer be working on the exact same requirements.

34 See, e.g., Broome (2013, esp. 153-54) and Kolodny (2005; 2007; 2008a; 2008b). See Daoust (2022, sect. 1) for a survey of putative structural requirements.

3.4. Escape route #3: Implicit Additional Conditions in the Account of Coherence

Another thing dispositionalists could reply is that there is an implicit condition of rational control in their account of coherence. Very roughly, rational control is the capacity to form and revise attitudes in accordance with the demands of rationality. Suppose the account of coherence is more explicit and includes a condition of rational control over the agent's attitudes, as in the following:

Coherence (With Rational Control). A set of attitudinal mental states is jointly incoherent iff it is (partially) constitutive of the mental states in question that, for any agent that holds these attitudes *and has rational control over them*, the agent is disposed, when conditions of full transparency are met, to give up at least one of the attitudes.

In the case of an Imperfect Optimizer, it seems that this condition is not satisfied, because Anna can't change some of her attitudes. So, the case of an Imperfect Optimizer would not be a counterexample to the dispositionalist account of coherence once it is more explicit.

Now, what is rational control, exactly? As I said earlier, we can make a distinction between structural and substantive rationality. Accordingly, here are two rough interpretations of rational control:[35]

1. *Structural Interpretation*. The degree of rational control agents have over their attitudes is determined by their capacity to form coherent combinations of attitudes.

2. *Substantive Interpretation*. The degree of rational control agents have over their attitudes is determined by their capacity to form attitudes that are supported by the reasons they have.

Begin with the Structural Interpretation. Under the Structural Interpretation of rational control, it appears that Anna has rational control over her attitudes. In all the cases described in section 2, one of the options she has is coherent. She can always make the conscious decision to form coherent combinations of attitudes. So, under the Structural Interpretation, Anna has rational control over her attitudes.

Let's now turn to the Substantive Interpretation. Suppose that, if we want to figure out whether Anna is coherent, we need to assume that she has the capacity to form attitudes that are supported by her reasons. This assumption would generate other problems for dispositionalists

35 See, for instance, Portmore (2019) for a much more precise and detailed formulation of the Substantive Interpretation.

like Worsnip. Here is a little bit of context: Worsnip (2018b) thinks that there are situations in which one's reasons decisively support incoherent combinations of attitudes. For instance, Worsnip thinks that one's reasons sometimes support incoherent combinations of beliefs, like "my evidence decisively supports believing P, but I do not believe P." Call this cases of Level-Splitting. Now, in accordance with his own account of coherence, Worsnip thinks that it would be difficult for agents to form such combinations of beliefs (under conditions of full transparency). So, in cases of Level-Splitting, transparent agents would have a limited capacity to form attitudes that are supported by their reasons. It would be difficult for them to form and sustain Inter-Level Incoherent combinations of attitudes, even if such attitudes are reasonable.

To make a long story short: For Worsnip, we can evaluate the coherence of agents without assuming that they have the capacity to form reasonable attitudes. In cases of Level-Splitting, Worsnip's account of coherence might even entail that agents like Anna lack the capacity to be maximally reasonable. And this is one of the reasons why people like Worsnip came up with accounts of coherence that are independent from substantive rationality.[36]

4. The Imperfect Optimizer and the Limits of First-Personal Deliberation

4.1. A Worry for Worsnip's Vindication of the Normativity of Structural Rationality

I just argued that the possibility of an Imperfect Optimizer raises a worry for the dispositionalist account of coherence. However, one could still think that we can define structural rationality in terms of the dispositions agents have. Perhaps we just need to revise the definition of coherence, so as to assume that agents *who do not face any mental constraints* would be disposed to be coherent. Or perhaps additional conditions in the definition of coherence (like a clause limiting coherence to "normal cases in which agents have control over all their attitudes") would allow us to get around the worry.[37] I leave it open whether we can unify coherence requirements based on the dispositions agents have in certain circumstances. However, I think that the possibility of an

36 Worsnip's (2018b) argument for the existence of conflicts between evidence-responsiveness and coherence requirements is, in part, in reaction to views that tie structural rationality to substantive rationality. For instance, according to some philosophers, one feature of structurally irrational attitudes is that they are guaranteed to be substantively irrational (Kiesewetter 2017; Kolodny 2007, 2008a; Lord 2018). Worsnip denies this. He says that, in cases of Level-Splitting, inter-level incoherent attitudes are structurally irrational, but not substantively irrational.

37 I raise another (tangential) worry against this account of coherence in Daoust (2022).

Imperfect Optimizer poses a more serious worry for Worsnip's answer to the Normative Question.

The possibility of an Imperfect Optimizer lends support against Worsnip's vindication of the normativity of structural rationality.[38] My reasoning goes as follows:

P1. Violating requirements like Probabilistic Consistency or Inter-Level Coherence is structurally irrational.

P2. For agents like the Imperfect Optimizer, some combinations of attitudes that violate these requirements can be the best option available to them (in terms of, e.g., expected accuracy or evidential support).

P3. Plausibly, agents like the Imperfect Optimizer should take the best option available to them.

P4. If agents like the Imperfect Optimizer treated some structurally irrational combinations of attitudes as off-limits in the course of first-personal deliberation, they would sometimes fail to take the best option available to them.

C. So, plausibly, the Imperfect Optimizer should *not* treat structurally irrational combinations of attitudes as off-limits in the course of first-personal deliberation.

I argued for P1 in section 3. I argued for P2 and P4 in section 2. C is supported by P1-P4. Also, C conflicts with Worsnip's vindication of the normativity of structural rationality. Recall that, for Worsnip, if some combinations of attitudes are structurally irrational, then agents should treat them as off-limits in first-personal deliberation. But it is hard to see why an Imperfect Optimizer should do that. Given our imperfections, we should sometimes consider the possibility of having incoherent combinations of attitudes.

What about P3? It is plausible to think that, epistemically speaking, agents like the Imperfect Optimizer should take the best available option they have. For instance, when rational agents deliberate on what to believe, they should try to achieve something epistemically valuable, like accuracy. And so, they should take the best option they have in terms of (expected) accuracy. Or agents should do their best to apportion their attitudes to their evidence. Accordingly, when rational agents deliberate on what to believe, they should take attitudes that best fit with their evidence. No matter how we interpret what's "best" exactly, P3 makes sense.

38 A related problem has been discussed in the practical realm—see Setiya's smoker (2007, 660) and Kiesewetter (2017, 99-102) for discussion. However, they discuss the implications of this kind of case for deontic detachment and transmission principles. I here discuss what we can learn from this case for the limits of first-personal deliberation (§4) and the separation of substantive and structural requirements (§5).

There are at least two possible responses to my argument. First, one could suggest that the argument merely entails this: agents like the Imperfect Optimizers fall under a *prima facie* obligation to be coherent, but this obligation can be overridden. Second, one could suggest that the conclusion of my argument does not matter, since we should merely care about the attitudes of perfect agents.

4.2. First Response: Imperfect Agents and Overridden Obligations

Begin with the first response. Could the Imperfect Optimizer have a *prima facie* but overridden obligation to be coherent? Worsnip seems to remain neutral on whether the obligation to be coherent can be overridden. At least, he doesn't mention this possibility. However, this response is compatible with his general framework. So, we can explore this response on Worsnip's behalf.

I do not have a conclusive argument against this possibility. However, I think that this possibility raises a small worry: while we usually still "feel the pull" of overridden *prima facie* obligations, it's hard to see the appeal of coherence in the situations experienced by Imperfect Optimizers. This suggests that we do not have an (overridden) obligation to be coherent. I also think that there is a more simple and natural interpretation of the case, in terms of nonseparable obligations.

First, the small worry: There is a common normative psychology associated with overridden obligations. That is, we usually still feel the pull of overridden obligations. Consider a common case of overridden obligation. Say, suppose Anna has a *prima facie* duty to keep her promises and a *prima facie* duty not to lie. Unfortunately, Anna can't fulfill both obligations simultaneously (say, she promises to Bob that she will tell Carla that P, and she later discovers that P is false). She has to make a choice between keeping her promise and not lying. Now, assume that her *prima facie* obligation to keep her promise is overridden. Yet Anna will still feel uneasy with her decision not to keep her promise. Her intention not to keep her promise leaves *normative residue*. Lying and not keeping our promises are sources of discomfort, and we normally feel compunction when we violate our *prima facie* obligations (Ross 2002, 28). In other words, there is a common normative psychology associated with the violation of *prima facie* obligations. We usually still feel the pull of obligations, even if they are overridden.

Now, consider the case discussed in section 2. Suppose Anna is imperfect and can't have the substantively ideal doxastic attitude towards P (or the ideal mental process for forming and revising attitudes). Anna can then see that, if she takes coherent combinations of attitudes, she just *aggravates* her own situation. That is, for imperfect agents like us, satisfying structural requirements of rationality will sometimes just *make things worse*. And so, it won't be difficult for Anna to consider incoherent combinations of attitudes in the course of her deliberation. It's hard to feel the pull of coherence in cases like this, precisely because being coherent would just make deliberation worse. Unless the putative obligation to be coherent has its own, unique normative psychology, this suggests that imperfect agents do not have an overridden obligation to be coherent.

Second, I think there is a more natural way to interpret the situation experienced by Imperfect Optimizers. Rather than saying that agents like Anna have an overridden obligation to be coherent, I want to suggest that we do not fall under a *separate* obligation to be coherent. The notion of (non)separability naturally accounts for the fact that imperfect agents like Anna are not required to be coherent. I will come back to this point in the conclusion.

4.3. Second Response: Imperfect Agents and the Norms of First-Personal Deliberation

Worsnip could argue that the conclusion of my argument does not matter. That is, he could claim that norms applying to non-ideal agents, like the Imperfect Optimizer, are irrelevant, and that we should care only about norms for perfect agents. So, perhaps the Imperfect Optimizer should not structure her deliberation in a coherent way, but this is irrelevant. What matters is how *perfect* agents structure their deliberation.

In response to this objection, note first that the Imperfect Optimizer fails to be perfect in the substantive sense. She does not respond correctly to her reasons. But she is not imperfect in the *structural* sense. She is in a position to satisfy all the requirements of structural rationality. So, there is still a sense in which the agent can be perfect: the Imperfect Optimizer can be fully *structurally* rational. Even if one merely cares about idealized norms, the argument has some import for perfect *structural* rationality.

One could respond that Worsnip's thesis is limited to agents that are in a position to be perfectly rational *both* in the structural sense and in the substantive sense. The Imperfect

Optimizer violates a substantive norm. So, the case is irrelevant. We can ignore the complications surrounding the Imperfect Optimizer.

I have two responses to this objection. First, recall that Worsnip's argument has to do with the first-personal deliberative perspective. Norms of good first-personal deliberation are generally sensitive to our imperfections. For instance, suppose I have conclusive reasons in favour of the following decisions:

1. Visit Grandma this Sunday.

2. Call Grandma in advance to tell her that I will visit.

Making these decisions is the "ideal" option. If I were perfect, this is what I would do. Unfortunately, I know that I am lazy, and that I will not visit her on Sunday. Yet I could still call Grandma and tell her that I will visit. What should I do? Given my imperfections, good first-personal deliberation will lead me not to call her. My first-personal deliberation will be sensitive to the fact that I am not perfect, and that I sometimes fail to do what my reasons support.[39] If I call her in advance to tell her that I will visit while knowing that I will not visit, I do more harm than good.

There are other ways in which norms of first-personal deliberation are sensitive to our imperfections. Just think of strategic reasoning. Suppose I am playing chess and I need to figure out my next move. I know how Deep Blue would figure out its next move. For IBM's computer, the best way to process the evidence is to do complex suppositional reasoning concerning billions of possibilities. However, I can't reason like this. Part of my imperfection is that I can't process billions of possibilities at once. My deliberation will be adapted to my limited cognitive capacities.[40] Once again, the upshot is that norms governing deliberation are sensitive to our imperfections. Accordingly, if we really care about norms of first-personal deliberation, we should not focus exclusively on what perfect agents would do and believe.[41]

39 See, e.g., Broome (2013, sec. 7.4) on similar problems.

40 Simon (1976).

41 Of course, you might not care about norms of first-personal deliberation. Perhaps you think that normativity has to do with a third-personal perspective. But then, this would raise a different worry for Worsnip and the possibility of vindicating the normativity of structural rationality.

Second, even if the case of imperfect agents fell outside the scope of Worsnip's thesis, it would still reveal that his vindication does not generalize well to a variety of cases. Specifically, it would be unclear why agents *like us* should care about structural rationality. Perhaps perfect cherubs should, in the course of first-personal deliberation, treat structurally irrational combinations of attitudes as off-limits. But why should agents *like us* care about structural rationality?

5. Conclusion and Discussion

I have presented a problematic case for Worsnip's answers to the Descriptive and the Normative questions. The Imperfect Optimizer violates some plausible substantive requirements of rationality. For instance, relative to her knowledge, some of her degrees of belief do not maximize expected accuracy. She could still satisfy plausible structural requirements of rationality, like Probabilistic Consistency. However, she has good reasons to violate these requirements: relative to the options she has, violating structural requirements might be the best option she has.

The case reveals that a vindication of the normativity of structural rationality is not independent of how well agents do in the *substantive* sense. If you fail to satisfy a substantive norm, it might be best for you to violate a structural norm as well. In technical terms: norms of structural and substantive rationality are not *separable*.[42] If we take this observation seriously, we should consider the possibility that a good vindication of the normativity of structural rationality is not independent from other norms and requirements.

The concept of nonseparability allows us to better grasp the permissions and obligations of Imperfect Optimizers. Nonseparable norms have certain logical properties. If my obligation to $(\varphi_1 \wedge \varphi_2)$ is nonseparable, then I do not fall under two separate obligations to φ_1 and to φ_2.[43] We see

42 See, e.g., DiPaolo (2019), Lipsey and Lancaster (1956) and Wiens (2020). Cases involving nonseparable variables can trigger second-best problems. Roughly, the idea is this: Suppose that the satisfaction of some conditions C_1, C_2,..., C_n characterize an ideal. If we can't satisfy C_i, then satisfying the remaining conditions might be suboptimal. For concreteness, suppose that the ideal of rationality is characterized by the satisfaction of structural and substantive requirements. However, suppose that an Imperfect Optimizer can't satisfy all the substantive requirements. Then, satisfying the remaining structural requirements might be suboptimal for her.

43 Formally: If $O(\varphi_1 \wedge \varphi_2)$ denotes a nonseparable obligation, then $O(\varphi_1 \wedge \varphi_2)$ does not entail $O(\varphi_1) \wedge O(\varphi_2)$. See Räikkä (2000).

this more clearly in cases of partial compliance with ideals. Suppose that I fall under the following obligation: Visit grandma this Sunday and call her in advance to tell her that I will visit. If I were the perfect grandchild, this is what I would do. The two parts of this obligation are in interaction and make sense *together*. To see this, suppose that a constraint refrains me from visiting Grandma this Sunday. Then, I should not call her, because telling her that I am visiting on Sunday would just aggravate the situation. In other words, the following reasoning is mistaken: "I won't fulfill my obligation to visit Grandma this Sunday, but the least I can do is satisfy my other obligations and tell her that I am coming." What I am required to do is the conjunction of two actions, not each individual action.

In light of these properties of nonseparable norms, we can provide a natural explanation of why imperfect agents like us do not have an obligation to be coherent in *abstraction from everything else*. As with the above case, substantive and structural norms interact with each other in important ways. These interactions suggest that the structural is not separable from the substantive. And so, if we take these observations seriously, we do not necessarily have a separate obligation to be coherent. This is especially clear in cases where we are less than ideally reasonable.

These lessons concerning nonseparability also matter for how we should address the normativity of coherence. For instance, Worsnip says that "coherence requirements and [norms of responsiveness to] evidential reasons must be separated and theorized separately" (Worsnip 2018b, 3). For him, this is a consequence of possible conflicts between the two. I think there are many different ways in which this claim makes sense. I agree, for instance, that the structural does not reduce to the substantive. And I agree that there are interesting phenomena that are proper to structural rationality. But we should not *always* address structural and substantive rationality separately from each other. We sometimes need to address them together. For instance, a good vindication of the normativity of rationality should not separate the structural from the substantive. There are normatively relevant interactions between these two spheres of rationality.

Here is an analogy with what I have in mind. Suppose you are a policy advisor. You can recommend policies that promote different values, like freedom, security, equality, and so forth. We care about all these values. But on occasion, given some constraints we face, we have to

make a trade-off between different values. For example, policies that contribute to security, like mandatory vaccination, come at the cost of freedom, and vice versa. This indicates that security and freedom are independent of each other, in the sense that we can't reduce one to the other. But we shouldn't conclude from this observation that we should address the issues of security and freedom separately from each other. In fact, if there are conflicts between different values in policy-making, this is a good reason to address them *together*, so as to strike the right balance between these values.

In the case of requirements of coherence, something similar is happening. There are interactions between norms of coherence and substantive norms. Specifically, for imperfect agents like us, coherence can come at the cost of making better decisions. This confirms that coherence and reasons-responsiveness are independent from each other, in the sense that we can't reduce one to the other. But surely, this isn't a reason to address them separately from each other all the time. If we want to strike the right balance between norms of coherence and norms of reasons-responsiveness, we need to address the two together, *as a whole*.

Finally, if we take the interactions between structural and substantive rationality seriously, we need to address several questions that can be divided in two groups:

1. **The nonseparability of substantive and structural requirements**: How do structural and substantive norms interact with each other? In this paper, we saw that they are not *separable*, in the sense that violating a substantive norm can be relevant for determining whether we should satisfy the structural norms (and vice versa). But in what sense are they nonseparable, exactly? There are many different ways in which norms can be nonseparable. Are they sometimes complementary? Are they sometimes synergistic? Do they stand or fall together? Can some structural requirements sometimes *compensate* for our violating substantive requirements (and vice versa)?
2. **The normativity of nonseparable requirements**: How do we vindicate the normativity of requirements that are not separable? The arguments for and against the normativity of structural rationality discussed in print tend to *isolate* structural rationality from everything else. That is, philosophers are trying to figure out whether agents should be structurally rational *independently of everything else*. However, under the assumption that structural

requirements are not separable from substantive requirements, is this the right way to vindicate the normativity of structural rationality? How can putative vindications of the normativity of structural rationality be sensitive to the nonseparability between structural and substantive requirements?

In a nutshell: we should care about the interactions between substantive and structural rationality. And this might require vindicating the normativity of rationality *as a whole*—that is, in ways that do not isolate the structural from the substantive.

Acknowledgments. This research was supported by the Fonds de recherche du Québec – Société et culture (grant #268137) and the Social Sciences and Humanities Research Council (grant #756-2019-0133). I warmly thank two anonymous referees from *Erkenntnis* for their constructive comments and genuinely interesting suggestions.

References

Arkes, Hal R., Gerd Gigerenzer, and Ralph Hertwig. 2016. "How Bad Is Incoherence?" *Decision* 3 (1): 20–39.

Broome, John. 2013. *Rationality through Reasoning*. Oxford: John Wiley & Sons.

Brunero, John. 2009. "Against Cognitivism about Practical Rationality." *Philosophical Studies* 146 (3): 311.

Brunero, 2020. *Instrumental Rationality: The Normativity of Means-Ends Coherence*. Oxford: Oxford University Press.

Daoust, Marc-Kevin. 2022. "Structural Rationality and the Property of Coherence." *Pacific Philosophical Quarterly* (online first): 1–25. DOI: 10.1111/papq.12404

Davidson, Donald. 1985. "Incoherence and Irrationality." *Dialectica* 39 (4): 345–354.

Davidson, Donald, John Charles Chenoweth McKinsey, and Patrick Suppes. 1955. "Outlines of a Formal Theory of Value, I." *Philosophy of Science* 22 (2): 140–160.

de Finetti, Bruno. 1937. "La Prévision: Ses Lois Logiques, Ses Sources Subjectives." In *Annales de l'institut Henri Poincaré*, 7:1–68.

de Finetti, Bruno. 1974. *Theory of Probability*. Wiley.

DiPaolo, Joshua. 2019. "Second Best Epistemology: Fallibility and Normativity." *Philosophical Studies* 176 (8): 2043-2066.

Dogramaci, Sinan. 2018. "Solving the Problem of Logical Omniscience." *Philosophical Issues* 28 (1): 107–128.

Easwaran, Kenny. 2011. "Bayesianism II: Applications and Criticisms." *Philosophy Compass* 6 (5): 321–332.

Easwaran, Kenny. 2016. "Dr. Truthlove or: How I learned to Stop Worrying and Love Bayesian Probabilities." *Noûs* 50 (4): 816-853.

Easwaran, Kenny, and Branden Fitelson. 2015. “Accuracy, Coherence, and Evidence.” *Oxford Studies in Epistemology* 5: 61–96.

Elga, Adam. 2010. “How to Disagree about How to Disagree.” In *Disagreement*, edited by Ted Warfield and Richard Feldman. Oxford: Oxford University Press.

Elga, Adam. 2013. “The Puzzle of the Unmarked Clock and the New Rational Reflection Principle.” *Philosophical Studies* 164 (1): 127–39.

Ekstrom, Laura Waddell. 1993. “A Coherence Theory of Autonomy.” *Philosophy and Phenomenological Research* 53 (3): 599–616.

Feldman, Richard. 2005. “Respecting the Evidence.” *Philosophical Perspectives* 19 (1): 95–119.

Fishburn, Peter. 1974. “Lexicographic Orders, Utilities and Decision Rules: A Survey.” *Management Science* 20 (11): 1442–1471.

Greaves, Hilary, and David Wallace. 2006. “Justifying Conditionalization: Conditionalization Maximizes Expected Epistemic Utility.” *Mind* 115 (459): 607–632.

Hájek, Alan. 2009. “Dutch Book Arguments.” In *The Handbook of Rational and Social Choice*, edited by Paul Anand, Prasanta Pattanaik, and Clemens Puppe. Oxford: Oxford University Press.

Horowitz, Sophie. 2014. “Immoderately Rational.” *Philosophical Studies* 167 (1): 41–56.

Horowitz, Sophie. 2019. “The Truth Problem for Permissivism.” *Journal of Philosophy* 116 (5): 237–262.

Jeffrey, Richard. 2004. *Subjective probability: The Real Thing*. Cambridge: Cambridge University Press.

Joyce, James M. 1998. “A Nonpragmatic Vindication of Probabilism.” *Philosophy of Science* 65 (4): 575–603.

Joyce, James M. 2009. “Accuracy and Coherence: Prospects for an Alethic Epistemology of Partial Belief.” In *Degrees of Belief*, edited by Franz Huber and Christoph Schmidt-Petri, 263–297. Dordrecht: Springer.

Kiesewetter, Benjamin. 2017. *The Normativity of Rationality*. Oxford: Oxford University Press.

Kolodny, Niko. 2005. “Why Be Rational?” *Mind* 114 (455): 509–563.

Kolodny, Niko. 2007. “How Does Coherence Matter?” *Proceedings of the Aristotelian Society* 107: 229–263.

Kolodny, Niko. 2008a. “The Myth of Practical Consistency.” *European Journal of Philosophy* 16 (3): 366–402.

Kolodny, Niko. 2008b. “Why Be Disposed to Be Coherent?” *Ethics* 118 (3): 437–463.

Lasonen-Aarnio, Maria. 2015. “New Rational Reflection and Internalism about Rationality.” *Oxford Studies in Epistemology* 5: 145–171.

Lee, Wooram. 2020. “The Real Myth of Coherence.” *Erkenntnis* (online first): 1-20. DOI: 10.1007/s10670-020-00239-y

Leitgeb, Hannes, and Richard Pettigrew. 2010a. “An Objective Justification of Bayesianism I: Measuring Inaccuracy*.” *Philosophy of Science* 77 (2): 201–235.

Leitgeb, Hannes, and Richard Pettigrew. 2010b. “An Objective Justification of Bayesianism II: The Consequences of Minimizing Inaccuracy*.” *Philosophy of Science* 77 (2): 236–272.

Lewis, C. I. 1946. *An Analysis of Knowledge and Valuation*. LaSalle (Ill): Open Court.

Lewis, David. 1971. “Immodest Inductive Methods.” *Philosophy of Science* 38 (1): 54–63.

Lipsey, Richard G., and Kelvin Lancaster. 1956. “The General Theory of Second Best.” *The Review of Economic Studies* 24 (1): 11–32.

Lord, Errol. 2014. “The Coherent and the Rational.” *Analytic Philosophy* 55 (2): 151–175.

Lord, Errol. 2018. *The Importance of Being Rational*. Oxford: Oxford University Press.

Meacham, Christopher JG. 2015. “Understanding Conditionalization.” *Canadian Journal of Philosophy* 45 (5–6): 767–797.

Meacham, Christopher JG. 2016. “Ur-Priors, Conditionalization, and Ur-Prior Conditionalization.” *Ergo, an Open Access Journal of Philosophy* 3 (17): 444-492. DOI: 10.3998/ergo.12405314.0003.017

Pettigrew, Richard. 2016. *Accuracy and the Laws of Credence*. Oxford: Oxford University Press.

Portmore, Douglas W. 2019. “Control, Attitudes, and Accountability” In *Oxford Studies in Agency and Responsibility 6*, edited by D. Shoemaker, 7-32. Oxford: Oxford University Press

Räikkä, Juha. 2000. “The Problem of the Second Best: Conceptual Issues.” *Utilitas* 12 (2): 204-218.

Raz, Joseph. 1992. “The Relevance of Coherence.” *BUL Rev.* 72: 273–321.

Ross, William David. 2002. *The Right and the Good*. Oxford: Oxford University Press.

Sayre-McCord, Geoffrey. 1985. “Coherence and Models for Moral Theorizing.” *Pacific Philosophical Quarterly* 66 (1–2): 170–190.

Schoenfield, Miriam. 2017. “Conditionalization Does Not (in General) Maximize Expected Accuracy.” *Mind* 126 (504): 1155–1187.

Setiya, Kieran. 2007. “Cognitivism about instrumental reason.” *Ethics* 117 (4): 649-673.

Simon, Herbert A. 1976. “From Substantive to Procedural Rationality”: 65-86. In *25 Years of Economic Theory*, edited by T. J. Kastelein, S. K. Kuipers, W. A. Nijenhuis, and G. R. Wagenaar. Boston (MA): Springer.

Staffel, Julia. 2019. *Unsettled Thoughts: A Theory of Degrees of Rationality*. Oxford: Oxford University Press.

Thagard, Paul. 2002. *Coherence in Thought and Action*. Cambridge (MA): MIT press.

Thorstad, David. 2021. “The Accuracy-Coherence Tradeoff in Cognition.” *British Journal for the Philosophy of science* (online first). DOI: 10.1086/716518

Tversky, Amos. 1969. “Intransitivity of Preferences.” *Psychological Review* 76 (1): 31–48.

Tversky, Amos, and Daniel Kahneman. 1983. “Extensional versus Intuitive Reasoning: The Conjunction Fallacy in Probability Judgment.” *Psychological Review* 90 (4): 293.

Wedgwood, Ralph. 2017. *The Value of Rationality*. Oxford: Oxford University Press.

Wiens, David. 2020. “The General Theory of Second Best Is More General Than You Think.” *Philosophers’ Imprint* 20 (5): 1-26. HDL: 2027/spo.3521354.0020.005

Worsnip, Alex. 2018a. “What Is (In)Coherence?” *Oxford Studies in Metaethics* 13: 184–20.

Worsnip, Alex. 2018b. “The Conflict of Evidence and Coherence.” *Philosophy and Phenomenological Research* 46 (1): 3–44.

Worsnip, Alex. 2019. “Disagreement as Interpersonal Incoherence.” *Res Philosophica* 96 (2): 245–268.

Worsnip, Alex. 2021. “Making Space for the Normativity of Coherence.” *Noûs* (online first), 1-23. DOI: 10.1111/nous.12362

Zynda, Lyle. 1996. “Coherence as an Ideal of Rationality.” *Synthese* 109 (2): 175–216.