

# Making Space for the Normativity of Coherence

Alex Worsnip

Forthcoming in *Noûs*

Penultimate version; please cite published version where possible

Some patterns of attitudes (and absences thereof) don't fit together right, in a distinctive sort of way: they are jointly *incoherent*. Examples include inconsistent beliefs, cyclical preferences, failures to intend the means to one's ends, and various forms of akrasia. Very plausibly, such incoherent combinations of states are irrational. (The particular kind of irrationality that they involve is often labelled *structural irrationality*.) Also very plausibly, facts about what is rational or irrational have normative significance – roughly, that is, they at least have some bearing on what attitudes we *ought* to have (or get ourselves to have).

And yet, as a raft of literature over the last two decades has made clear, there are deep puzzles about how facts about (in)coherence could be normatively significant. The deepest such puzzle, in my view at least, is what I'll call the problem of *making space* for the normativity of coherence. The problem, in rough outline, is this: if considerations of coherence are normatively significant, then it seems that we ought to take account of them in our deliberation and reasoning. But it is unclear exactly what role they should play in our deliberation and reasoning. Considerations of coherence don't seem to show up in reasoning about what to believe, intend, desire, hope, fear, and so on; moreover, they seem awkward to take account of *alongside* more “substantive” considerations about the merits of such attitudes. This casts doubt on the normative significance of considerations of coherence.

I'll try to make this problem more precise in due course. I'll then try to solve it. My view will make more sense a bit further into the dialectic, but for those who don't like to be kept in suspense, the slogan is this: considerations of coherence constitute reasons for *structuring deliberation in certain ways*; more particularly, in ways that treat incoherent combinations of attitudes as off-limits, and that focus one's deliberation on adjudicating between the coherent combinations. Not only does this account address the challenge of making space for the normativity of coherence; it also constitutes an original, positive account of *how* coherence is normatively significant. Despite the voluminous literature on the “normativity of rationality,” very few such accounts have been forthcoming in the extant literature.

## 1. “The normativity of rationality”: getting clearer on the question

In this preliminary section, I will situate the question I'm interested in by locating it within the broader debate about the “normativity of rationality”. To outsiders, this debate may seem puzzling. Of course, it might be thought, claims about rationality are normative (rather than, say, descriptive) claims. Yet prominent philosophers profess to be skeptics about the normativity of rationality,<sup>1</sup> while others devote whole books to defending it.<sup>2</sup> To understand how there is a question of substance to be

---

<sup>1</sup> Most famously, Niko Kolodny (2005), and on one interpretation, Joseph Raz (2005). In his book *Rationality Through Reasoning*, John Broome also takes such skepticism very seriously, admitting that he has no argument to rule it out (see Broome 2013: 204-5).

<sup>2</sup> E.g. Kiesewetter (2017), Wedgwood (2017), and Lord (2018).

investigated here, we need to get clearer both about what is the notion of rationality, and about the notion of normativity, that is in play.

### 1.1 *What kind of rationality is at issue?*

As others have argued, there are two different notions of rationality that need to be distinguished: substantive, and structural.<sup>3</sup> *Substantive* rationality, roughly, consists in having attitudes (and absences thereof<sup>4</sup>) – and, perhaps, performing actions – that are *reasonable* or *justified*. On a prominent and plausible account,<sup>5</sup> this amounts to their being supported by (and based on) one’s evidence-relative reasons, where the qualifier ‘evidence-relative’ functions to exclude any reasons of which one is justifiably ignorant given one’s evidence. When we say that it’s irrational for an educated person in 2020 to believe that climate change is a hoax, or that it’s irrational to (intentionally) leave one’s umbrella at home when it’s clearly about to rain, these are most naturally read as ascriptions of substantive irrationality. In these cases, the agent is failing to respond correctly to her (evidence-relative) reasons against for belief and (intentional) action.

*Structural* rationality, by contrast, consists roughly in having attitudes that *cohere* with one another. On a prominent and plausible account,<sup>6</sup> this amounts to their respecting various constraints that forbid certain clashing combinations of attitudes: for example, constraints that forbid inconsistent beliefs, means-end incoherence, cyclical preferences, various forms of akrasia, and so on. When we say that it’s irrational to have the combinations of attitudes that violate these constraints – for example, that it’s irrational to intend to go to the bar given your intention to do everything you can to avoid getting COVID-19 and your belief that to do that, you need to refrain from going to the bar, or that it’s irrational to believe the plane will crash when *you yourself acknowledge* that you have no evidence that it will crash, these are most naturally read as ascriptions of structural irrationality. In these cases, the agent’s attitudes are incoherent with one another.<sup>7</sup>

In this paper, I want to focus on the normativity of *structural* rationality, or coherence. My own view is that there is in fact no interesting debate to be had about the normativity of *substantive* rationality, but I do not need to rely on that here: I will just stipulate that it is structural rationality that I am interested in, and understand that qualifier to be implicitly in place in what follows.<sup>8</sup>

---

<sup>3</sup> For this distinction, see esp. Worsnip (2018b, forthcoming), Fogal (2020), and Fogal & Worsnip (forthcoming). Others who have used the term “structural rationality” in at least roughly the same way include Scanlon (2007), Chang (2013), Wallace (2014), Neta (2015), and Kiesewetter (2017).

<sup>4</sup> In what follows, I’ll omit the qualifier ‘and absences thereof’ for brevity, using ‘attitudes’ more broadly as including absences of (positive) attitudes.

<sup>5</sup> Cf., *inter alia*, Kiesewetter (2017) and – though his view is not quite the same – Lord (2018).

<sup>6</sup> Cf., *inter alia*, Broome (2013).

<sup>7</sup> Some have challenged the fundamental distinctness of substantive and structural rationality, and/or have thought that structural rationality is just about responsiveness to “subjective” reasons. For a defense of the distinction, and of the view that structural rationality cannot be understood in terms of responsiveness to *any* kind of reasons, see Fogal & Worsnip (forthcoming); Worsnip (forthcoming: chs. 1-4).

<sup>8</sup> In making this stipulation, I am following some of the most prominent early contributors to the debate on the normativity of rationality (Kolodny (2005: 509-10; Southwood 2008: 9-10). Some more recent contributors to the debate (Kiesewetter 2017; Lord 2018) have broken with this earlier approach. They think, instead, that the question ‘is rationality normative?’ can be sensibly asked without any prior disambiguation of ‘rationality’, and then answered by conducting a substantive, non-stipulative investigation into what rationality, *simpliciter*, consists in: either reasons-responsiveness *or* coherence. And

Notably, despite the voluminous literature on the normativity of rationality, defenders of the normativity of *structural* rationality, or coherence, are rare,<sup>9</sup> and accounts of *how* structural rationality is normative are rarer still. In time, I aim to provide just such an account.

## 1.2 What is it for rationality to be *normative*?

What do participants in the debate mean when they ask whether rationality is *normative*? Most participants in the debate cash out what it would be for rationality to be normative in terms of the notion of a *reason*. This is in many respects a natural approach. Normative claims are often glossed as claims that are concerned with what we *ought* to do, believe, intend, desire, and so on. But some consideration can be normatively significant without always tracking what one all-things-considered ought to do (or believe, intend, desire, etc.) – since considerations of genuine normative significance can be outweighed. Reasons to  $\Phi$  are often understood as the considerations that together add up to determine whether you ought to  $\Phi$ . By bearing on what you ought to do (or what attitudes you ought to have), reasons are the considerations that have normative significance. So the notion of a reason is a natural one to look to in glossing what it is for something to be normative.

However, there are different ways to use the notion of a reason to cash out the thesis that rationality is normative. One way one sometimes seemed the thesis cashed out, either explicitly or implicitly, is like so:<sup>10</sup>

**Reason(s) to be Rational in General.** There is at least one fact F such that F is a reason for all agents to have all of the responses that are rationally required of them.

However, I don't think this is a good way to cash out the thesis that rationality is normative. Given this claim, the natural next question is, "...and what *is* that reason?"<sup>11</sup> Candidate answers, at this point, are ones like the following: being rational will help us to achieve our ends; being rational will help us to do what we "objectively" ought to do;<sup>12</sup> being rational is necessary for self-governance;<sup>13</sup> being rational stops us from being vulnerable to Dutch books and money pumps; being rational makes us interpretable to others in a social environment.<sup>14</sup>

---

they think that if they can show that rationality consists in reasons-responsiveness, and *not* coherence, then they have thereby shown that it is normative. In my view, this approach to the debate is mistaken, because the notions of rationality as reasons-responsiveness (substantive rationality) and that of rationality as coherence (structural rationality) do not compete as rival accounts of rationality *simpliciter*. That is, no showdown is required as to whether rationality *is* one or *is* the other; they are simply two different deontic notions, both of which have their own distinctive roles to play, and both of which can be picked out with the ordinary word 'rationality'. Consequently, there is no substantive project of figuring out which of the two rationality "is", and there is no question of whether rationality *simpliciter* is normative that could be answered by executing that project.

<sup>9</sup> Again, both Kiesewetter (2017) and Lord (2018) argue for the normativity of rationality, but only via identifying rationality *simpliciter* with substantive rationality. They are skeptics about the normativity of structural rationality.

<sup>10</sup> Cf., e.g., Kolodny (2005); also, e.g., Mildemberger (2019).

<sup>11</sup> Cf., e.g., Kolodny (2005: 542).

<sup>12</sup> Cf. Wedgwood (2017).

<sup>13</sup> Cf. Bratman (2009).

<sup>14</sup> Cf. Mildemberger (2019).

Note, however, that there is something striking about all of these possible answers. All of them make our reasons to have the rational responses essentially *derivative* on something else: our reasons to achieve our ends, our reasons to do what we objectively ought to do, our reasons to be agents, or our reasons to be interpretable to others in a social environment (or, perhaps, the *value* of these various things). I think that the move toward these derivative pictures is the inevitable result of posing the question in such a way that what is demanded is a reason to be rational in general. The very question forces us to look outside of rationality itself to find some external goal or value on which its significance is derivative. I do not think that someone who thinks that rationality is normative should automatically be forced into this model, where rationality's normative significance is derivative on some non-rationality-constituted (e.g. moral, prudential, etc.) consideration.<sup>15</sup> Hence, I think, Reason(s) to be Rational in General is a bad gloss on the normativity of rationality.

What can we replace it with? A better way to understand the thesis that rationality is normative, I suggest, is something like this:

**Facts about Rationality Constitute Reasons.** Any fact about what rationality requires of some agent A constitutes a reason for A to have some appropriately related response R.<sup>16</sup>

Here, it is the fact about rationality itself that *constitutes* the reason to have a particular response. This much better accommodates those who think that rationality's normative significance is non-derivative. What they can say is that facts about rationality constitute reasons, whether or not there is some (non-rationality-constituted, e.g. prudential) reason to be rational in general. This view is perfectly intelligible, as we can see by analogy with those who hold analogous thesis about the non-derivative normative significance of other sorts of facts. For example, moralists who think that morality is non-derivatively normatively significant think that moral considerations<sup>17</sup> constitute reasons whether or not there is some (non-moral, e.g. prudential) "reason to be moral" in general.<sup>18</sup> Similarly, Humeans who think that our desires are non-derivatively normatively significant think that facts about what we desire constitute reasons whether or not there is some (non-desire-given, e.g., prudential) "reason to satisfy our desires" in general.<sup>19</sup> Of course, any one of these views can be questioned. But plausibly, on pain of infinite regress, *everyone* is going to have to hold that *something* is of non-derivative normative significance, where this significance isn't explained in terms of some further reason of a different kind.<sup>20</sup>

At the same time, Facts about Rationality Constitute Reasons also accommodates those who think that facts about rationality only *derivatively* constitute reasons – i.e. those who think that facts

---

<sup>15</sup> Cf. Southwood (2008: esp. 18).

<sup>16</sup> This can be thought of as a precisification of the somewhat metaphorical claim that rationality "gives" us reasons (cf. Broome 2005a).

<sup>17</sup> The phrase 'moral considerations' is deliberately neutral as to whether it's the fact *that* some action is morally required that constitutes the reason or the fact(s) that *make it the case that* the action is morally required that constitute the reason. Dancy (2000: 165-7) and Zimmerman (2007: Appendix 2) hold that it's the latter. But Johnson King (2019) persuasively argues that *either* fact can correctly be cited as a reason (even though the two reasons don't aggregate).

<sup>18</sup> Cf. Prichard (1912). Southwood (2008: 17) also point out this analogy.

<sup>19</sup> Cf. Schroeder (2007: 212-4, 217).

<sup>20</sup> Cf. Korsgaard (1996: 33).

about rationality constitute reasons, but only in virtue of some more basic (prudential, moral, etc.) “reason to be rational” in general – as a proponent of the normativity of rationality. It’s thus in itself neutral on whether rationality is derivatively or non-derivatively normatively significant. This makes it a more ecumenical formulation of the thesis that rationality is normative.

A few further remarks about Facts about Rationality Constitute Reasons. First, the term “appropriately related” is of course somewhat imprecise. On the simplest model, the appropriately related response will just be the one that is rationally required of A. Thus, any fact of the form *rationality requires A to  $\Phi$*  will be a reason for A to  $\Phi$ . However, I intend “appropriately related responses” to also cover some other responses. For example, suppose that facts of the form *rationality requires A to  $\Phi$*  are reasons for A to do something that *results* in her  $\Phi$ -ing. This would count as vindicating Facts about Rationality Constitute Reasons as I’ve formulated it. But not *all* potential responses are appropriately related. For example, suppose that facts of the form *rationality requires A to  $\Phi$*  are reasons to mock A if she  $\Phi$ ’s. I take it this would not constitute a vindication of the normativity of rationality in the sense that advocates of that thesis are after. So I am not counting mocking A if she  $\Phi$ ’s as an appropriately related response. Though the boundaries as to what counts as an appropriately related response are of course fuzzy, I rely on the reader to have some sense on what counts as a closely enough related response to count as a vindication of the normativity of rationality in an interesting sense.

Second, in introducing the distinction between substantive and structural rationality above, I said that structural rationality, in contrast to substantive rationality, is not about responding to reasons. But if Facts about Rationality Constitute Reasons is true (with ‘rationality’ being understood as picking out structural rationality), then there are special reasons *of* structural rationality. Are these two claims consistent? Yes. Here is an analogy. Suppose, for the sake of argument, that considerations of beauty constitute reasons: the beauty of some X constitutes a reason to bring X about. Were this true, it would not follow that what it is to be beautiful is to be responsive to reasons, or even to be responsive to the beauty-reasons (i.e. the reasons constituted by considerations of beauty). After all, beauty can be a property of things that cannot respond to reasons. And even restricting ourselves to things that can respond to reasons – humans, say – it’s possible to be beautiful without having made oneself so in response to a beauty-reason. Thus (now exiting the scope of our supposition for the sake of argument), though our account of what beauty is will not invoke reasons-responsiveness, we can nevertheless go on to ask whether facts about beauty constitute reasons. Similarly, even if our account of structural rationality does not invoke reasons-responsiveness, we can go on to ask whether facts about structural rationality constitute reasons.

Finally, since it is *structural* rationality we are concerned with here, and structural rationality consists in coherence, we can capture the central question at issue here either by asking whether facts about *rationality* (and irrationality) constitute reasons or by asking whether facts about *coherence* (and incoherence) constitute reasons. I will treat these two formulations interchangeably and move between them in what follows.<sup>21</sup> Given the framing in terms of coherence, we can also use the term ‘coherence-

---

<sup>21</sup> It might be contended that although the property of structural (ir)rationality and the property of (in)coherence are coextensive, they are not strictly speaking the same property. In particular, it may be that they have an asymmetrical metaphysical relationship, where facts about structural (ir)rationality obtain *in virtue of* facts about (in)coherence, not the other way round. Someone who thinks this might go on to contend that strictly speaking, it’s the facts about (in)coherence *rather* than the facts about structural (ir)rationality that constitute reasons for appropriately related

based reasons’ to refer to reasons constituted by facts about the (in)coherence of attitudes. Our question, in a nutshell, is whether there are such reasons (for appropriately related responses).

## 2. The problem of making space for the normativity of coherence

I’ll now introduce what I consider to be the deepest challenge for the normativity of structural rationality, or coherence. I’ll call this challenge, which is inspired by Kolodny (2005),<sup>22</sup> the problem of *making space* for the normativity of coherence. Broadly speaking, the challenge is to explain how considerations of coherence could fit into our deliberation or reasoning about which attitudes to adopt.

To spell the problem out, it’s helpful to focus on an example. Let’s take the requirement of “inter-level coherence”, which requires coherence between one’s first-order and higher-order beliefs: more particularly, it requires (among other things) that one believe  $p$  if one believes that one’s evidence decisively supports believing  $p$ .<sup>23</sup>

Here’s the core of the challenge. If considerations of coherence constituted reasons, then when I believed that my evidence decisively supports believing  $p$ , it would seem to me that I have *two* reasons to believe  $p$ . First, the fact that my evidence decisively supports believing  $p$  (or whatever fact or facts make it the case that it decisively supports believing  $p$ ). And second, the fact that, since I *believe* that my evidence decisively supports believing  $p$ , believing  $p$  would make me coherent (or, bring me to satisfy a requirement of structural rationality). But that is not, in fact, how things seem in situations like this one. In situations like this one, I won’t treat this coherence-constituted consideration as a second reason to believe  $p$ , in addition to that provided by my evidence. From the first-personal point of view, coherence-based reasons seem to be “superfluous”; they don’t “add anything”.

There are several different, albeit related, points here. The first is that the above illustration brings out how considerations of coherence seem typically not to, indeed perhaps can’t, enter our deliberation about what attitudes to have. When I deliberate about whether to believe  $p$ , I just think directly about my evidence for  $p$ . I don’t think about whether believing  $p$  would make me coherent. And there seems nothing misplaced about this. Yet, on a popular view, reasons to  $\Phi$  are the sorts of things that can (and, when everything is going right, do) show up in deliberation about whether to  $\Phi$ . This casts doubt on whether considerations of coherence genuinely constitute reasons.

The second point is that if I did, in my deliberation, treat the fact that believing  $p$  would make me coherent as a second, additional reason to believe  $p$ , I would seem to be making a mistake. At least assuming that I’m right that my evidence does decisively support believing  $p$ , to say that I have two

---

responses (just as some hold, as we saw in fn. 17, that it’s the fact(s) *in virtue of* which one is morally required to  $\Phi$ , rather than the fact *that* one is morally required to  $\Phi$ , that constitute one’s moral reason to  $\Phi$ ). As with the moral case, my own view, in line with Johnson King (2019), is that either framing of the reason is permissible, even assuming that the two properties are not identical. But someone who thinks that only the framing of the reason in terms of coherence is strictly correct should be able to accept the full substance of my positive account, and if they do, they will also count as a proponent of the normativity of coherence (and in a good sense, I think, the normativity of structural rationality).

<sup>22</sup> More specifically, it is the *third* of three challenges for the normativity of structural rationality that Kolodny explores. I find the first two less powerful.

<sup>23</sup> The label ‘inter-level coherence’ is due to Worsnip (2018a). Kolodny calls (this dimension of) inter-level coherence “B+” (Kolodny 2005: 521).

reasons to believe  $p$  – first, the fact that my evidence decisively supports believing  $p$ , and second, the fact that, since I *believe* that my evidence decisively supports believing  $p$ , believing  $p$  would make me coherent – seems to be engaging in an illicit kind of “double-counting”. Again, assuming that the first, evidential consideration *is* a reason to believe  $p$ , this suggests that the second, coherence consideration is not a second, additional reason to believe  $p$ .

The third point is this. It’s generally thought that only evidence for  $p$  can constitute the “right kind of reason” (or, as I’ll write for short, a “right-kind” reason) for belief for  $p$ . Anything that isn’t evidence can at best constitute a “wrong-kind reason” for belief. But the fact that believing  $p$  will make me coherent is not (usually) in and of itself evidence for  $p$ . So the fact that believing  $p$  will make me coherent seems to be, at best, a wrong-kind reason for belief.

According to some philosophers,<sup>24</sup> so-called “wrong-kind reasons” aren’t, in fact, genuine normative reasons at all. (A popular argument for this view appeals to the claim that wrong-kind reasons can’t feature in our deliberation, or motivate us, thus bringing the first and third points together.) If that’s so, it would follow straightaway from coherence considerations being wrong-kind that they aren’t reasons. However, even for those who do believe in wrong-kind reasons, there is a challenge here. For (as I’ll discuss further in §6 below), the paradigmatic wrong-kind reasons for believing  $p$  are those that show believing  $p$  to be *valuable* (for example, cases where one is offered a large sum of money to believe  $p$ ). But it isn’t clear that it *is* (especially) *valuable* to have coherent attitudes, at least not intrinsically or in all cases. So coherence considerations seem to fit *neither* the mold of right-kind reasons, nor the mold of wrong-kind reasons, for attitudes. If that is so, then again, one may suspect that they are no kind of reason at all.

In the foregoing, I’ve explained the problem of making space for the normativity of coherence in terms of the example of how considerations of coherence seem deliberately superfluous as reasons for *belief*, when added to *evidential* considerations. But all three aspects of the challenge generalize to other attitudes, and to the superfluosity of coherence-based reasons on top of the relevant “right-kind” reasons. For example, consider intention. It seems that I do not treat the fact that intending to  $\Phi$  would make me coherent (say, because I believe I ought to  $\Phi$ ) as an additional reason to intend to  $\Phi$  on top of the substantive (moral, prudential, etc.) reasons that, in my view, make it the case that I ought to  $\Phi$ . And if I did, I would be engaged in a form of illicit double-counting. Moreover, on the standard view, some reason to intend to  $\Phi$  is right-kind only if it is also a reason *to*  $\Phi$ . But the fact that it would be incoherent to fail to *intend* to  $\Phi$  doesn’t seem to be a reason to  $\Phi$ . Thus, again, it seems at best a wrong-kind reason to intend to  $\Phi$ .

### 3. Toward a solution: the form of coherence requirements

We are considering the problem of how coherence considerations could enter into deliberation about what to believe, intend, etc. But there is a tacit assumption in the way that the challenge was set up. The assumption was that the if coherence considerations fit into this deliberation, they would do so by appearing as considerations that favor particular attitudes. For example, the assumption was, the

---

<sup>24</sup> Including Kolodny himself (Kolodny 2005: 548-51). See also, e.g., Shah (2006); Parfit (2011: Appendix A); Way (2012); Rowland (2014); Kiesewetter (2017: 11-13).

fact that believing  $p$  would make me coherent (i.e., bring me to satisfy a requirement of structural rationality), given my other background mental states, would appear in my deliberation as a consideration in favor of believing  $p$ . As we have just seen, this model seems fairly unpromising. If this is the only way that coherence consideration could enter into deliberation, we might conclude that coherence considerations cannot play any kind of role in deliberation. But I want to explore whether there's a different way of fitting coherence considerations into deliberation.

Let's work up to this by considering the form of coherence requirements. On my view,<sup>25</sup> coherence requirements can be understood, fundamentally, as prohibitions on combinations of attitudes.<sup>26</sup> (I assume that prohibitions can still be called 'requirements', since a prohibition on a combination is equivalent to a requirement not to have that combination.) So, for example, the inter-level coherence requirement prohibits you from having either of the following combinations: {believing that your evidence decisively supports some doxastic attitude  $D$ , lacking  $D$ }, {believing that your evidence does not adequately support some doxastic attitude  $D$ , having  $D$ }. The instrumental requirement (roughly) prohibits you from {intending to  $\Phi$ , believing that  $\Psi$ -ing is necessary for  $\Phi$ -ing, failing to intend to  $\Psi$ }. The noncontradiction requirement prohibits you from {believing  $p$ , believing not- $p$ }. The transitivity requirement prohibits you from {preferring  $A$  to  $B$ , preferring  $B$  to  $C$ , preferring  $C$  to  $A$ }. And so on. To be structurally rational, on my view, is to satisfy all of these coherence requirements.

This account is broadly in the spirit of what is often called a "wide-scope" account of structural rationality, most prominently associated with John Broome.<sup>27</sup> On the wide-scope view, rational requirements can be stated using a deontic operator that takes wide-scope over a material conditional. So, for example, the instrumental requirement can be stated as follows: rationality requires of you that ((you intend to  $\Phi$  & you believe that  $\Psi$ -ing is necessary for  $\Phi$ -ing)  $\rightarrow$  you intend to  $\Psi$ ). Depending on whether one accepts some non-trivial assumptions,<sup>28</sup> this may in fact be equivalent to my formulation (though I prefer mine for presentational reasons).<sup>29</sup> In any case, the crucial point is that on both of these ways of stating coherence requirements, coherence requirements generally do not tell you to have, or not to have, some *individual* attitude.<sup>30</sup> Call any view of coherence requirements with this feature 'wide-scope in spirit'. For example, suppose that you intend to  $\Phi$ , believe that  $\Psi$ -ing is necessary for  $\Phi$ -ing, but do not intend to  $\Psi$ . On a wide-scope in spirit view, the instrumental requirement is silent on how to resolve this situation: whether to do so by coming to intend to  $\Psi$ , by giving up your intention to  $\Phi$ , or by giving up your belief that  $\Phi$ -ing is necessary for  $\Psi$ -ing.<sup>31</sup> Of

---

<sup>25</sup> Defended at more length in Worsnip (forthcoming: ch. 6).

<sup>26</sup> I do not mean to say that all prohibitions on combinations of attitudes are coherence requirements.

<sup>27</sup> See Broome (1999, 2013: ch. 8) for classic statements of this view. Forerunners of Broome's view include, among others, Greenspan (1975) and Dancy (1977).

<sup>28</sup> Namely, that one can substitute logical equivalents within the scope of deontic operators, and that a requirement *that* it's not the case that you  $\Phi$  is equivalent to a requirement *not to*  $\Phi$ .

<sup>29</sup> Again, see Worsnip (forthcoming: ch. 6).

<sup>30</sup> The exception is the degenerate case when a single attitude on its own is incoherent, and is thereby prohibited by a coherence requirement. The usual example is a single belief in the conjunctive proposition ( $p$  & not- $p$ ).

<sup>31</sup> Schroeder (2004) and Kolodny (2005) thought this was a flaw: this is their "symmetry objection". For responses, see Way (2011) and Brunero (2012). Incidentally, as a matter of dialectic, I take it that Kolodny's third challenge to the normativity of coherence requirements – the one that I'm interested in here – is supposed to stand independently of his contention that such requirements are narrow-scope. It's only his first challenge that is supposed to rest on that premise.



course, that's not to say that there aren't further normative requirements that determine a unique way to go among these options; it's just to say that by the lights of the instrumental requirement *itself*, any of these ways of going is OK. All it says is that you'd better go one of these ways; it's not permitted to remain instrumentally incoherent.

Given this, we might ask: if coherence requirements don't (generally) *require* individual attitudes of us, then why should the proposal that they are normative be taken to mean that they generate *reasons* for individual attitudes, either?

Now, it isn't *obvious* that prohibitions on combinations of attitudes can't generate reasons for individual attitudes. Here is one potential model for how they might. Suppose again that you believe that your evidence decisively supports believing  $p$ . Now, we might say, one thing that speaks in favor of believing  $p$  is that believing  $p$  would bring you into satisfaction of a coherence requirement. This is an instance of what we can call the 'satisfaction-as-reason model':

**Satisfaction-As-Reason Model.** The fact that  $\Phi$ -ing would bring you into satisfaction of a coherence requirement, given your other states, is a reason to  $\Phi$ .

To say that  $\Phi$ -ing would bring you into satisfaction of a coherence requirement isn't to say that  $\Phi$ -ing is the *only* way of bringing you into satisfaction of a coherence requirement. Thus, the satisfaction-as-reason model is compatible with recognizing the point that coherence requirements are prohibitions on combinations of states, and don't generally mandate individual attitudes.

Still, it doesn't seem that someone who thinks that coherence requirements prohibit combinations of attitudes is *forced* to adopt the satisfaction-as-reason model. This would only follow given two ancillary assumptions: first, that you have a reason to satisfy the coherence requirement; and second, that reasons "transmit" in a strong sense, whereby if you have a reason to  $\Psi$ , then for *any*  $\Phi$  that results in your  $\Psi$ -ing, you have a reason to  $\Phi$ .<sup>32</sup> Such transmission principles are dubious, at least in full generality; I'll return to this point at the end of §6. Provided that we're not forced to adopt the satisfaction-as-reason model, thinking of coherence requirements as requiring the avoidance of certain combinations of attitudes – rather than as requiring individual ones – opens up the way for alternative models for the normativity of coherence whereby coherence considerations count in favor of something other than individual attitudes. And since the satisfaction-as-reason model does not offer any solution to Kolodny's initial challenge about how coherence considerations, qua reasons for individual attitudes, could fit into deliberation, it's worth pursuing such alternative models.

Strangely, though, this seems to be about where much of the existing literature gives out. For example, Broome defends a wide-scope account of coherence requirements, and he's also at least sympathetic to the view that such requirements are normative. But he cannot accept the satisfaction-as-reason model, for it would lead to "bootstrapping" of a kind that he considers unacceptable.<sup>33</sup> If the satisfaction-as-reason model were correct (in full generality), then whenever I intend to  $\Phi$  and I believe that  $\Psi$ -ing is necessary for  $\Phi$ -ing, I'd thereby have a reason to intend to  $\Psi$  (since doing so

---

<sup>32</sup> Raz (2005) is sometimes read as relying on something like this principle. See Broome (2005b) and Kiesewetter (2017: 92-95) for discussion.

<sup>33</sup> See Broome (2005b).

would bring me into satisfaction of a coherence requirement). Moreover, it would seem that I could make it the case that I have this reason by adopting the relevant intention and belief – no matter how dastardly or worthless the intended end, and no matter how crazy or evidentially-unsupported the means-end belief. Broome thinks that this kind of “bootstrapping” shouldn’t be possible, and that this is what rules out “narrow-scope” views of rational requirements.<sup>34</sup> But if the satisfaction-as-reason model were correct, it would also affect Broome’s own wide-scope view.<sup>35</sup> Thus, Broome can’t accept the satisfaction-as-reason model. And yet Broome doesn’t ever explore what the alternative model for the normativity of coherence considerations might be: what they might be reasons *for*, if not for individual attitudes. I want to explore just such an alternative model. To set the stage, I’m going to make some remarks about deliberation.

#### 4. Deliberation as holistic

Here’s a pretty obvious fact that philosophers nevertheless sometimes fail to attend to sufficiently: when we deliberate, we often don’t deliberate about individual attitudes in isolation. Suppose I’d like to wear some new shoes to the wedding I’m attending on Saturday, instead of my old, beat-up pair. But suppose also that the only way I’ve thought of for getting hold of new shoes involves paying a visit to the megamall to purchase some. And suppose I really don’t want to pay a visit to the megamall. There are three attitudes to settle: first, whether to intend to wear new shoes to the wedding, second, whether to believe that in order to wear new shoes to the wedding, I need to visit the megamall; and third, whether to intend to visit the megamall. But it would be odd if I deliberated about each of these attitudes one by one, in isolation. Instead, it’s more natural to consider them together. I’ll engage in a more holistic deliberation process, thinking: what’s more important to me, being able to wear the new shoes to the wedding, or avoiding a visit to the megamall? Moreover, if both of those are looking really important to me, I might well think: is there some feasible way that I could get hold of some new shoes *without* visiting the megamall?

Moreover, in these deliberative contexts, I can settle on more than one attitude simultaneously. If I’ve fully settled on the belief that visiting the megamall is necessary for getting hold of the shoes, then I can then deliberate about whether to visit the megamall or give up on the idea of getting new shoes for the wedding. While I’m still deliberating, I’m in an unresolved state with respect to both decisions, but when I decide, I can resolve both simultaneously: for example, by settling on the intention to wear new shoes to the wedding and the intention to visit the megamall simultaneously. I don’t have to start with the former and then *reason* to the latter.

Similarly, this simultaneous settling can happen with a belief and an intention. Suppose I decide that I care about getting the shoes above everything else in play, and settle on that intention. I can then deliberate about whether it’s really necessary to go to the megamall in order to get the shoes, and I can conclude that deliberation simultaneously with the belief that it is necessary and the intention to go to the megamall.

---

<sup>34</sup> See, e.g., Broome (2004: 30, 49).

<sup>35</sup> Cf., again, Kiesewetter (2017: 92-95).

I've focused here on the case where I'm initially undecided about all three of the relevant options. But similar points hold when I have positive attitudes at the start of my deliberation. Suppose I've been intending, for months, to wear new shoes to the wedding, but I also have a standing intention never to visit the megamall. Suddenly I realize: the wedding is *this* Saturday, and it's too late for an online order! To get new shoes, I now believe, I'm going to have to visit the megamall. At this point, my intentions are up for *revision* simultaneously. My decision problem is this: either I've got to abandon my intention to wear new shoes to the wedding, or I've got to abandon my intention not to visit the megamall (and, indeed, come to positively intend to visit it), or I've got to find adequate reason to doubt that I really need to visit the megamall in order to get hold of the shoes. Again, it's natural for me to deliberate about each of these possible *revisions* of my attitudes in a holistic manner, rather than in isolation.

Notice that for all this to work in the way I've described, however, I need to be treating the combination of {intending to wear new shoes to the wedding, believing that in order to wear new shoes to the wedding I have to visit the megamall, not intending to go to the megamall} as off-limits. I may not have – indeed, probably haven't – done this by explicitly considering this combination of attitudes and consciously ruling it out. Nevertheless, I don't treat it as a live option in my deliberation. By doing this, I structure my deliberative problem, so that my options become *either* abandoning my intention to wear new shoes to the wedding, or coming to intend to visit the megamall, or giving up my belief that I need to go to the megamall to get new shoes. Moreover, once I've deliberated in such a way as to further narrow down this field of options to two – for example, by settling on the belief that in order to wear new shoes to the wedding I have to visit the megamall – my deliberations about the other two attitudes are tied together. Since the only options I'm still taking seriously are {intending to wear new shoes, intending to visit the megamall} and {not intending to wear new shoes, not intending to visit the megamall}, I make my decision about whether to adopt both intentions simultaneously.

Similar points hold for deliberation in the theoretical realm. Suppose I'm trying to figure out what to think about, say, whether a Democrat will be US President in 2025. Typically, I don't consider two separate questions in isolation: first, whether (to believe that) the evidence supports believing that a Democrat will be US President in 2025; and second, whether (to believe that) a Democrat will be US President in 2025. Nor do I even deliberate about the former question, form a belief about it (e.g.: my evidence does support believing that a Democrat will be US President in 2025), and then engage in some *reasoning* to get to the belief that a Democrat will be US President in 2025. Rather, I tie my deliberation about these two questions together. Again, I can do that by treating certain combinations of attitudes as off-limits. Specifically, letting  $d$  be the proposition that a Democratic will be US President in 2025, I treat both {believing that the evidence supports believing  $d$ , not believing  $d$ }, as and {believing that the evidence doesn't support believing  $d$ , believing  $d$ }, as off-limits. This means that I am effectively deliberating only between two options: {believing that the evidence supports believing  $d$ , believing  $d$ }, and {not believing that the evidence supports believing  $d$ , not believing  $d$ }. So I make up my mind about both questions together simultaneously and in concert. Again, by treating certain combinations of attitudes as off-limits, I structure my deliberative problem, tying together two questions in thought.

But now notice two things. First, in both of the examples just considered, the combinations of attitudes that I treat as off-limits are precisely the combinations that are incoherent: those that are means-end incoherent, and inter-level incoherent, respectively. And second, when I treat these combinations of attitudes as off-limits, and deliberate between the remaining, coherent combinations of attitudes, it's then the case that whichever attitudes I ultimately settle on in my deliberation, I will end up satisfying the relevant coherence requirement (in this instance). So, whether I settle on not intending to wear new shoes to the wedding, on not believing that visiting the mall is necessary for doing this, or on intending to visit the megamall, I'll end up satisfying the means-end coherence requirement. Similarly, whether I settle on {believing that my evidence supports *d*, believing *d*} or on {not believing that my evidence supports believing *d*, not believing *d*}, I'll end up satisfying the inter-level coherence requirement. At least as long as my attitudes are formed through deliberation, I will avoid the incoherent combinations of attitudes that I am treating as off-limits.

## 5. Coherence considerations as reasons to structure deliberation

So far I've been describing a way in which it's possible for me to treat certain (incoherent) combinations of attitudes as off-limits in my deliberation, thus structuring my deliberative problem, and tying together my deliberations about multiple attitudes. Moreover, in doing so, I make it the case that the attitudes that result from my deliberation will respect coherence requirements.

I am not, of course, saying that agents always *do* treat these incoherent combinations of attitudes as off-limits. However, I now want to make the following normative claim: the ways of deliberating that I described in the last section match how we *should* deliberate. That is: we *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.

Moreover, I claim, the fact that some possible combination of attitudes is incoherent (or structurally irrational) is the reason that one should treat it as off-limits in deliberation. Since, in treating certain attitudes as off-limits (and focusing deliberative attention on others), one thereby structures one's deliberation in a certain way, we can say for short that considerations of coherence are reasons to *structure deliberation* in certain ways. This lands us at my proposal:

**Reasons-to-Structure-Deliberation Model.** Considerations of coherence constitute reasons to structure deliberation in certain ways. More specifically: the fact that some possible combination of attitudes is incoherent is a reason to treat it as off-limits in one's deliberation.

Whenever I talk about structuring deliberation in what follows, I mean that as shorthand for the kind of structuring of deliberation that is spelled out here: namely, treating certain possible combinations

of attitudes as off-limits, and focusing one's deliberative attention on others.<sup>36</sup> What is it to treat a combination of attitudes as off-limits in deliberation? Negatively, it involves simply not considering it (or being disposed not to consider it) as a deliberative option – not taking it seriously as a combination of attitudes one might adopt, not entertaining the substantive merits or demerits of adopting it. Somewhat more positively, however, it also involves considering (or being disposed to consider<sup>37</sup>) various different ways of *avoiding* the relevant incoherent combination, and to deliberate *between* those ways of avoiding the relevant incoherent combination, weighing their substantive merits and demerits.<sup>38</sup> I take it, then, that treating a combination of attitudes as off-limits involves certain mental performances or dispositions toward such performances; it is not itself an attitude, such as belief.<sup>39</sup>

As previously noted, treating an attitude as off-limits will also place certain constraints on the outcomes of one's deliberation: that if one adopts one attitude, one thereby adopts (or avoids) another as well. For example, if one is treating the incoherent combination of {believing  $p$ , believing not- $p$ } as off-limits, one ties together believing  $p$  with refraining from believing not- $p$ , so that the latter comes along with the former (and likewise for believing not- $p$  and refraining from believing  $p$ ). Similarly, if one treats the incoherent combination of {intending to  $\Phi$ , believing that to  $\Phi$  one must  $\Psi$ , not intending to  $\Psi$ } as off-limits, then once one has settled on the belief that to  $\Phi$  one must  $\Psi$ , one ties together intending to  $\Phi$  with intending to  $\Psi$ , so that the latter comes along with the former (and likewise for not intending to  $\Psi$  and not intending to  $\Phi$ ).

We could also perhaps get at what I'm describing with the phrase 'treat a combination as off-limits' by saying that someone who does this has effectively ruled the incoherent combinations of attitudes out in advance of substantive deliberation about which attitudes to have. But this language is liable to mislead. It may make it sound as if a good deliberator should start deliberation (or some pre-deliberative phase) by thinking of all the (relevant) incoherent combinations of attitudes and explicitly ruling them out. This is rarely what happens, and would be an absurd demand. Rather, a good deliberator is usually predisposed to display implicit sensitivity to the incoherence of certain combinations, and to treat those combinations as off-limits automatically, without any need to explicitly consider them and rule them out. This is enough to count as successfully responding to her coherence-based reasons to treat incoherent combinations as off-limits.

That said, some incoherences can be subtler than others, and sometimes, a deliberator may find herself explicitly entertaining a combination of attitudes that is, as a matter of fact, incoherent. On my view, in these situations, the deliberator has a coherence-based reason – constituted by the incoherence of the combination in question – to *explicitly* rule that combination of attitudes out from

---

<sup>36</sup> Likely, there are other things that also qualify as ways of structuring deliberation in a perfectly good sense, such as choosing to tackle some complex decision problem in some specific sequence. There will be other reasons bearing on these other ways of structuring deliberation.

<sup>37</sup> In cases where one is deliberating in the first place, where it doesn't just seem *obvious* how to settle the relevant deliberative question, etc.

<sup>38</sup> Is treating a combination of attitudes as off-limits (i.e., "structuring deliberation" in the particular way I'm interested in) itself a *part* of deliberation? This depends on how widely we construe 'deliberation': if it is restricted to the process of weighing the pros and cons of the options one is taking seriously, then structuring deliberation is distinct from it; but if it includes the matter of which options one considers or takes seriously in the first place, then it's a part of it, albeit a special part.

<sup>39</sup> Consequently, it doesn't itself enter into coherence relations, which hold between attitudes (and absences thereof).

her deliberation. Moreover, in such cases, one has reason to rule out this combination *immediately*, on the grounds that it is incoherent, without first engaging in (or finishing up) deliberation about the merits of the individual attitudes that constitute the combination.

What emerges from all this is a picture whereby coherence-based reasons, on one hand, and “substantive” (viz. moral, prudential, epistemic, etc.) reasons, on the other hand, play different sorts of roles in the deliberative process. The coherence-based reasons are reasons to structure the deliberative process in a certain way: they bear on which attitudes one should take seriously in one’s deliberation. Conversely, the substantive reasons enable one to adjudicate the merits of the options that one does take seriously, within the deliberative constraints set by the coherence-based reasons. Once my deliberation has been focused on the coherent options, my deliberation between those options should appeal neither to the facts about what my existing attitudes are, as such (though it could appeal to the contents of some of those attitudes) or to facts about coherence. After all, considerations of coherence by definition can’t help to adjudicate between coherent options. At that point, it’s substantive considerations all the way home.

To illustrate, return to our central example. Suppose I’ve responded to my coherence-based reasons by structuring my deliberation to focus on the options of *either* refraining from believing that going to the megamall is necessary for wearing new shoes to the wedding, *or* refraining from intending to wear new shoes to the wedding, *or* coming to intend to visit the megamall. Since all three of these attitudes suffice for avoiding means-end incoherence, considerations of coherence by definition can’t help me adjudicate between these three options. Rather, I should settle which of these three options to take by considering the substantive merits of each attitude (or absence thereof): the strength of my evidence that going to the megamall is necessary for wearing new shoes to the wedding, and the relative strength of the practical reasons for intending to acquire new shoes for the wedding versus those against intending to take a trip to the megamall.

## 6. How this solves the problem

Hopefully, it is now becoming clear how the reasons-to-structure-deliberation model addresses the problem of making space for the normativity of coherence. Recall that the challenge had three facets. The first was that considerations of coherence seem not to enter our deliberation about which attitudes to have, and this casts doubt on their status as reasons. The reasons-to-structure-deliberation model locates a way for them to play a role in deliberation. The role they play, however, is not in weighing alongside moral, prudential, and epistemic reasons in adjudicating the merits of individual attitudes, but rather in structuring our deliberation so as to treat incoherent combinations of attitudes as off-limits and focus deliberative attention on the remaining options.

This also helps us to see how the second facet of Kolodny’s challenge, which considered “double-counting”, is dispelled. Since coherence considerations, according to the reasons-to-structure-deliberation model, are not reasons *for adopting particular attitudes*, they don’t get counted alongside the substantive reasons for particular attitudes, and so there is no issue of double-counting in the offing.

Indeed, the reasons-to-structure-deliberation model can actually help to *explain* why it would be weird for coherence considerations to show up in deliberation as considerations in favor of individual attitudes. Go back to Kolodny’s example, and suppose that you believe that your evidence decisively supports  $p$ . Kolodny thinks that it would be weird to treat the fact that you believe this, and that believing  $p$  would therefore make you coherent, as a reason to believe  $p$ . With our account in hand, we can now say that if you’re even in this position in the first place, something has gone wrong in your deliberation. For in this situation, by hypothesis, you already believe that your evidence decisively supports  $p$ , but you’re still deliberating about whether to believe  $p$ . But this is already to have failed to structure your deliberation in a way such that you treat the option of {believing that your evidence decisively supports  $p$ , not believing  $p$ } as off-limits. If you’re deliberating about whether to believe  $p$  now, you’re deliberating between believing  $p$  and not believing  $p$ . If you’re doing this against the background of being settled in thinking that your evidence decisively supports  $p$ , then you seem to be taking seriously the option of {believing that your evidence decisively supports  $p$ , not believing  $p$ }. And your deliberation should have been structured so as to treat this option as off-limits.

The third facet of Kolodny’s challenge concerned the charge that coherence-based reasons would, if they were reasons at all, be “wrong-kind reasons”. Kolodny is not alone in assuming this.<sup>40</sup> Even those who are friendly to coherence-based reasons – such as John Broome (2013: 83-4) – concede that such reasons would have to be wrong-kind reasons.<sup>41</sup> The usual thought behind this, as we already saw, is that the only right-kind reasons for belief are evidential reasons. So if coherence considerations constitute reasons for belief, they must be wrong-kind reasons. (And similarly for other attitudes, the right-kind reasons for each of which must fit a specific description that coherence considerations always seem not to fit.)

But if coherence considerations are not reasons for attitudes at all, but rather for structuring deliberation in certain ways, this argument does not go through. For while right-kind reasons for *belief* may be exhausted by evidential considerations, this doesn’t mean that right-kind reasons for *structuring deliberation* are exhausted by (or even include) evidential considerations. Thus, considerations of coherence, construed as reasons to structure deliberation, need not be wrong-kind reasons.

Indeed, I think there is a good case to be made that they are right-kind reasons.<sup>42</sup> To do this, I need to first say what the right/wrong-kind distinction comes to. Following Howard (2019), I think the best way to draw the distinction is to identify right-kind reasons with “fit-related” reasons and

---

<sup>40</sup> In addition to Kolodny and Broome, see e.g. Kiesewetter (2017: 96-8, 106-8). One *apparent* exception to this trend is Jonathan Way (2010), who defends the view that when some means is necessary for an end, this very fact constitutes a *right-kind* reason to either intend the means or not intend the end. However, this reason is not itself constituted by a fact about coherence, so it isn’t a coherence-based reason in my sense. Indeed, like everyone else, Way actually seems to share the assumption that coherence of coherence *constituted* reasons, they would be wrong-kind reasons (*ibid.*: 219).

<sup>41</sup> All of the authors referenced in the previous footnote use the terminology of “object-given” and “state-given” reasons. I avoid this terminology because I don’t find the object/state language perspicuous, and because insofar as I think it makes sense, it doesn’t perfectly track the distinction that is fundamental here (see also Schroeder 2012).

<sup>42</sup> A further possibility is that they are neither right-kind nor wrong-kind reasons, because structuring deliberation is a kind of (mental) action. Some contend that the right-kind/wrong-kind distinction does not apply to actions (Heuer 2011: 179-80; Hieronymi 2013: 118; but see Schroeder 2010 for dissent). I am not persuaded that the right-kind/wrong-kind distinction *never* applies to actions: as I’ll go on to suggest, if the relevant action can be understood as a kind of response to an object that it is either fitting or unfitting given the properties of the object, then I think the right-kind/wrong-kind distinction can be made to apply. But even if actions do not admit of the right-kind/wrong-kind distinction, the crucial point for blocking Kolodny’s argument is that coherence considerations are *not* wrong-kind reasons.

wrong-kind reasons with “value-related” reasons. On this view, roughly, a right-kind reason to  $\Phi$  is a consideration that bears on the fittingness of  $\Phi$ -ing, whereas a wrong-kind reason to  $\Phi$  is a reason that bears on the value of  $\Phi$ -ing. So, for example, evidence for  $p$  bears on the fittingness of believing  $p$ , and so is a fit-related (right-kind) reason to believe  $p$ , whereas an offer of money to believe  $p$  bears on the value of believing  $p$ , and so is a value-related (wrong-kind) reason to believe  $p$ . Similarly, consider reasons for fear. Suppose there’s a big, aggressive tiger in front of you, that is gnashing its teeth and preparing to pounce. These scary features of the tiger make it fitting to be afraid of it, and are thus fit-related (right-kind) reasons to be afraid. Yet suppose also that tigers can smell fear, and that your being afraid makes it far more likely that the tiger will attack you. This fact makes it disvaluable to be afraid, and is thus a value-related (wrong-kind) reason not to be afraid.

With the distinction drawn in this way, it seems plausible that – once we understand coherence considerations as constituting reasons to structure deliberation in particular ways – they are right-kind reasons. For it is fitting to structure deliberation in ways that respect coherence constraints. This is true on either of two (non-identical) ways of understanding the term ‘fitting’. On the first (broader) way, to say that it’s fitting to  $\Phi$  is roughly to say that it’s  $\Phi$ -ing *correct* to  $\Phi$ , given the standards constitutive of  $\Phi$ -ing. And it is plausible that it is constitutive of deliberation that it’s correct to treat incoherent combinations of attitudes – like {intending to  $\Phi$ , believing that  $\Psi$ -ing is necessary to  $\Phi$ , not intending to  $\Psi$ } – as off-limits in one’s deliberation, and to focus one’s deliberation on deciding between coherent combinations of attitudes. To the extent that one treats incoherent combinations of attitudes as live in one’s deliberation, one seems to be engaging in deliberation that is incorrect of faulty *qua* deliberation.

On the second (narrower) way of understanding ‘fitting’, fittingness requires more structure, in that we can only apply fittingness-talk to responses to objects: some response  $R$  towards an object  $O$  is fitting just if  $O$  *merits* or *calls for*  $R$ . For example, the tiger *merits* being feared, and a proposition can *merit* or *be worthy of* belief. It might initially seem that a mental performance like structuring deliberation doesn’t have this additional structure. But consider what kind of structuring of deliberation is recommended (by coherence-based reasons) on my view: ruling out incoherent combinations, and focusing deliberation on coherent combinations. I think it’s true, in just the same sense that a proposition can *merit* being believed, that an incoherent combination of attitudes *merits* being treated as off-limits (in virtue of its incoherence). Thus, it’s fitting, even in the second sense, to treat it as off-limits.

There are also some factors that support identifying coherence considerations (construed as reasons to structure deliberation in certain ways) as right-kind reasons that hold independently of identifying the right/wrong-kind distinction with the fit/value-related distinction. First, it seems that we can directly and consciously respond to coherence considerations in the way we structure our deliberation: when I take some combination of attitudes to be incoherent, I can treat it as off-limits in my deliberation, explicitly on the basis that it is incoherent. This is typical of right-kind, but not wrong-kind, reasons,<sup>43</sup> the latter of which famously present difficulties in being responded to directly

---

<sup>43</sup> You might think that these considerations are easy to respond to just because structuring deliberation is an action, and all reasons for actions are easy to respond to. But this isn’t right: see the next footnote.



and consciously: I cannot believe that the number of the stars in the sky is even, or fear a butterfly, directly and consciously on the basis of your offering me money to do so.

Second, and relatedly, we can dream up certain wrong-kind reasons to structure deliberation, and when we do so, that sets up a conflict of reasons in which coherence-based reasons play the same role that other right-kind reasons play in analogous conflicts. So, for example, if I offer you money to treat an incoherent combination of attitudes as live in your deliberation, this feels very analogous to the just-mentioned cases where I offer you money to believe that the number of stars in the sky is even, or to fear a butterfly.<sup>44</sup> The analogy is not just that in both cases you're being offered money to do something (in the broadest sense of 'do'): it goes deeper than that. In both cases, you're being offered money to do something that doesn't really make sense on its own terms. You are, I suggest, being given a wrong-kind reason to go against what your right-kind reasons support.<sup>45</sup>

To finish this section, I want to briefly explain how understanding coherence considerations as (right-kind) reasons for structuring deliberation helps us to resist the aforementioned satisfaction-as-reason model, and the bootstrapping concerns that attend it. Recall that I said in §3 that someone who thinks that coherence requirements prohibit combinations of attitudes need only accept the satisfaction-as-reason model given two ancillary assumptions: first, that you have a reason to satisfy the coherence requirement; and second, that reasons "transmit" in a strong sense, whereby if you have a reason to  $\Psi$ , then for *any*  $\Phi$  that results in your  $\Psi$ -ing, you have a reason to  $\Phi$ . The reasons-to-structure-deliberation model makes both assumptions dubious. First, considerations of coherence in the first instance constitute reasons to structure one's deliberation in certain ways, rather than directly to satisfy coherence requirements. And second, considerations of coherence are right-kind reasons, which plausibly don't transmit in the requisite way.<sup>46</sup> For example, suppose you have an evidential reason to believe  $p$ , and due to some psychological quirk believing  $q$  will result in your believing  $p$ . It doesn't follow that you have an evidential (or other right-kind) reason to believe  $q$ . Thus, if coherence considerations are right-kind reasons, we shouldn't expect them to transmit this easily, either.<sup>47</sup>

---

<sup>44</sup> Notice that the offer of money to treat an incoherent option as live is, like other wrong-kind reasons to go against one's right-kind reasons, plausibly quite difficult to directly and consciously respond to. Certainly one could *think about* the incoherent option, and maybe even its pros and cons, in response to pragmatic reasons. But it seems to me that this doesn't on its own suffice for treating the option as live: to count as treating it as live one has to be genuinely considering *whether* to adopt the incoherent combination, where that involves *taking it seriously* or having some kind of genuine *openness* to adopting it in a way that does not seem possible in response to pragmatic reasons.

<sup>45</sup> Perhaps the wrong-kind reason is strong enough to make it the case that you ought, all-things-considered, to treat the incoherent combination of attitudes as live in your deliberation in this case. I take no stand on this; presumably, we'll want to generalize whatever we say about conflicts between right-kind and wrong-kind reasons more generally to this case. However, it does seem to me plausible that coherence considerations exhaust the *right-kind* reasons for treating combinations of attitudes as off-limits or live in deliberation. If this is so, there's at least a good sense of 'ought' – the one indexed only to the right-kind reasons – in which you don't just have *reasons* to structure deliberation in coherent ways, but indeed *ought* to do so. This may help to explain the sense that the normativity of coherence is "stringent" (cf., e.g., Way 2010: 228-9).

<sup>46</sup> Cf. Way (2010: 226-7); Kiesewetter (2017: 95-96). It's odd that Kiesewetter says that a general transmission principle for reasons is plausible (*ibid.*: 93), while then going on to admit that it's not plausible for right-kind reasons (*ibid.*: 96), since he elsewhere commits himself to the view that there are no wrong-kind reasons (*ibid.*: 12).

<sup>47</sup> By contrast, wrong-kind reasons seem to transmit more easily. If there are reasons that show  $\Psi$ -ing to be *valuable*, then any  $\Phi$  that results in your  $\Psi$ -ing will realize that value, and that seems to be a value-related reason to  $\Phi$ .

## 7. Two questions addressed

### 7.1 *Is this really a vindication of the normativity of coherence?*

The key feature of the reasons-to-structure-deliberation model that enables it to meet the challenge of making space for the normativity of coherence is that it reimagines what coherence-based reasons are reasons *for*. Whereas Kolodny assumed that coherence-based reasons would have to be reasons for individual attitudes, the reasons-to-structure-deliberation model proposes that they are instead reasons for structuring deliberation in certain ways – namely, treating certain combinations of attitudes as off-limits, and focusing one’s deliberative attention on the coherent options. But if my model says that considerations of coherence aren’t (directly) reasons *to have coherent attitudes*, does my model really vindicate the normativity of coherence?

The answer is that it does, if we take the definition of the normativity of coherence that I arrived at in §1. Recall that, as I defined it, coherence (or structural rationality) is normative if facts about coherence (or structural rationality) constitute reasons for some *appropriately related* response. Moreover, while the phrase ‘appropriately related’ is vague, I explicitly said that doing something that *results* in one’s being coherent counts as an appropriately related response. But structuring deliberation so as to treat incoherent combinations as off-limits, and to focus deliberation on the coherent options, *is* doing something that results in one’s being coherent. When done successfully, it ensures that, whichever attitudes one concludes deliberations with, these attitudes will be coherent. Thus, the reasons-to-structure-deliberation model is one on which coherence is normative in the sense that I defined.

You may think that this is all a cheat; that I have cooked the books by defining the normativity of coherence so as to count my own view as securing it. All I can say to this charge is that it does seem to me that a view on which (i) considerations of coherence have normative significance as reasons; and (ii) correctly responding to those reasons will result in one’s being coherent, counts as vindicating the normativity of coherence in a strong, robust, interesting sense. You may insist on understanding the claim that coherence is normative as restricted to the view that considerations of coherence directly constitute reasons for the individual attitudes that would make you coherent. If you define it that way, then I deny the thesis, since – as this section has made clear – I don’t think this is the way in which considerations of coherence are normatively significant. I am at peace with that. But, as my view illustrates, it is a large – and illicit – leap to conclude from this that coherence has no normative significance.

One further point here. My model is consistent with saying that you *ought* to satisfy the coherence requirements – in some sense of ‘ought’ that is, if perhaps not itself all-things-considered, at least *bears* on the all-things-considered ‘ought’. It’s just that the *way* that you should satisfy them is via their structuring your deliberation, not via your treating coherence considerations as reasons in favor of or against individual attitudes. This, I hope, should further allay the concern that my view isn’t one on which coherence is normative in a robust sense.

## 7.2 *Are coherence-based reasons still superfluous in a good sense?*

Behind Kolodny's challenge was the thought that if considerations of coherence constituted reasons, they would be superfluous. As I cashed this thought out, the idea was that, assuming that coherence considerations are supposed to be reasons for individual attitudes, they seem redundant when deliberating about which individual attitudes to have. This precise version of the challenge has been met by construing coherence considerations as reasons for structuring deliberation in certain ways, which locates a role for coherence considerations in deliberation. But there's a slightly different version of the superfluousness challenge that might be thought to still apply even to my view.

The thought is this. According to several philosophers,<sup>48</sup> you can never be in a position such that your substantive (i.e., moral, prudential, epistemic, etc.) reasons permit an incoherent set of attitudes. But if that's so, then it might seem *unnecessary* to structure your deliberation by treating incoherent sets of attitudes as off-limits. If you correctly deliberate about individual attitudes on their substantive merits, you'll never end up with incoherent attitudes anyway. In that sense, coherence-based reasons might still be thought to be superfluous.

As a matter of fact, I reject the claim that your substantive reasons never permit an incoherent set of attitudes.<sup>49</sup> This claim fails *either* if there are permissive cases, in which one's reasons merely *permit* attitudes that are jointly structurally irrational, *or* if there are conflict cases, in which one's reasons actually *require* attitudes that are jointly structurally irrational.<sup>50</sup>

But let's suppose, for the sake of argument, that the claim your substantive reasons never permit an incoherent set of attitudes is correct. If this is so, then it's true that coherence requirements don't impose any additional demands *on your attitudes*, over and above those imposed by your substantive reasons. Nevertheless, the present account isolates a *further* role that coherence considerations should play in your deliberation. Suppose you took incoherent combinations of attitudes seriously in your deliberation, and did not focus deliberation solely on the coherent combinations. Even if you were to then go on to rule out each of the incoherent combinations on the grounds that they contain attitudes for which you lack sufficient substantive reasons, and thus ultimately end up with coherent attitudes, there would be something defective about your deliberation *qua* deliberation due to your having taken the incoherent combinations seriously to start with. You shouldn't have to rule out the incoherent combinations on substantive grounds; their very incoherence

---

<sup>48</sup> Cf. Kolodny (2007); Kiesewetter (2017: ch. 9); Lord (2018: ch. 2).

<sup>49</sup> For arguments against this claim, see Worsnip (2018a, forthcoming: ch. 3).

<sup>50</sup> Countenancing conflict cases might open up a different worry, which is that it won't be appropriate to treat an incoherent combination of attitudes as off-limits in such cases, since doing so will lead one into a failure of reasons-responsiveness. (Thanks to an anonymous referee for raising this challenge.) This challenge deserves a fuller reply than I can give here (see Worsnip forthcoming: §8.8.4 for more), but briefly: the cases of conflict cases that I think are most plausible are ones where one has misleading higher-order evidence about what one ought to believe or do (cf. Worsnip 2018a). However, it's a feature of such cases that precisely because one's evidence is *misleading*, one isn't in a position to know (or justifiably believe) that it's misleading, and thus isn't in a position to recognize that one is in a conflict case. As a result, from one's own perspective, there will be nothing to rationalize treating the incoherent combination as a live option in this case. Consequently, I think it is still inappropriate to treat the incoherent combination of attitudes as a live option even in these cases. At least as far as *deliberating well* goes, it's still appropriate to treat the incoherent combination of attitudes as off-limits. The tragedy of conflict cases is that doing so will lead one into a failure to respond to one's reasons, but such is life if there are conflicts of this kind.

should be enough. The point here is not just the instrumentalist one that the incoherence of a combination of attitudes can give you a kind of “shortcut” in deliberation, one that helps you to focus your deliberation more efficiently on the options that are worth taking seriously – though this is also true. It is that there is something defective, something incorrect, about deliberation that doesn’t take this “shortcut”, that waits for the incoherent combinations to be ruled out through substantive consideration of their merits.<sup>51</sup>

## 8. Coda

I hope to have offered an attractive solution to the problem of making space for the normativity of coherence – namely, that coherence considerations are right-kind reasons to structure deliberation in certain ways. But some will feel that the problem I’ve addressed, and the solution I’ve offered, don’t speak to the heart of their angst. They will want to press a further question that they take to be thus far unanswered: “but *why* structure deliberation in the ways that coherence recommends?”

But in light of the considerations I surveyed in §1, we shouldn’t expect this question to have an informative answer; at least not if we think that considerations of coherence *non-derivatively* constitute reasons. In the same way as the moralist won’t think there’s a general, non-moral reason to do as morality requires, and the Humean won’t think there’s a general, non-desire-constituted reason to satisfy your desires, someone who thinks that considerations of coherence *non-derivatively* constitute reasons won’t think that there is a general, non-coherence-based reason to structure your deliberation in coherent ways. The best we can do is to give the rather uninformative answer “because it’s fitting to do so”.

It’s not fully clear what those who are unsatisfied with this are looking for, though. I can only speculate that they are looking for some explanation of why it’s *valuable* to structure one’s deliberation in coherent ways. But I do not think the normativity of coherence should be staked on such claims about value. Recall that in §6, I argued that once they are understood as reasons to structure deliberation, coherence-based reasons can be understood as *right-kind* reasons, where right-kind reasons are understood in terms of *fittingness*, not value. Return to the analogy of the scary tiger who can smell fear. As the tiger is gnashing its teeth and preparing to pounce, you have strong reasons to be afraid: these are right-kind, fit-related reasons. To deny that that there are such reasons is, to my

---

<sup>51</sup> I’m assuming here that one must either treat incoherent combinations of attitudes as off-limits in one’s deliberation (never even considering their substantive merits), or treat them as live (ruling them out, if at all, only on the basis of substantive considerations). An anonymous referee asks whether this assumption is correct. Could there be someone who counts as doing neither of these things – and who still ends up responding correctly to all her substantive reasons, without treating incoherent combinations of attitudes as off-limits *or* ruling them out on substantive grounds? I can imagine only two candidates for a person who might answer to this description. A first is someone who never treats incoherent attitudes as live in her deliberation, but not in any way that shows even implicit sensitivity to their incoherence. But this person seems to avoid treating incoherent attitudes as a live option only by a massively unlikely accident of luck. Since we cannot reasonably expect ourselves to ourselves to be the beneficiaries of this kind of extreme luck, this seems to me not really to show that coherence consideration are redundant for us in deliberation. The second candidate is someone who just never thinks about combinations of attitudes together, always deliberating about them one by one in isolation. But this would be a *really* bizarre person. For example, she would have to deliberate *separately* about whether (to believe) it’s raining, and whether (to believe) it’s not raining, instead of just thinking about whether it’s raining, and simultaneously settling on believing it’s raining and not believing that it’s not raining. It does seem to me that this would be faulty deliberation.

mind, is to be willfully unresponsive to our ordinary practices of normative thought and talk. However, given that being afraid will make the tiger much more likely to attack you, it might not be at all *valuable* to be afraid. If there's any value to having fitting fear-states *as such* (and even this seems dubious), it is infinitesimal compared with the value of getting attacked by a tiger. Thus, the (strong) fit-related reasons to be afraid in this case just aren't staked on the *value* of being afraid. Similarly, I suggest, the (strong) fit-related reasons to structure one's deliberation in coherent ways just aren't staked on the *value* of doing so.

Some philosophers may dispute this, because they're skeptical that there are *any* reasons that don't bottom out in realizing or promoting value.<sup>52</sup> Such philosophers would have to deny that you do have reasons to fear the tiger, or reasons to believe what your evidence supports, independently of doing so realizing or promoting value. To this, I have two responses, one defensive and one attacking. The defensive response is that even staying neutral on these claims, I think it is a major achievement if we can assimilate skepticism about the normativity of coherence to skepticism about the normativity of right-kind, or fit-related, reasons more generally. The thesis that there are coherence-based reasons has typically been taken to pose special problems that the thesis that there are fit-related reasons for fear, or (fundamentally) evidential reasons for belief, has not. (Indeed, some of the major skeptics about the normativity of coherence, far from being skeptics about these right-kind reasons, are skeptics about *wrong-kind* reasons.<sup>53</sup>) It's a major step forward if we can get these various putative reasons onto equal terms.

The attacking response – though I can offer it as nothing more than a polemic here – is that we should resist this value-centric way of thinking. The considerations that (rightly) get a grip on us in the space of reasons are not just those considerations that concern our realizing value. They are also those considerations attention to which enables things – the external world, our minds, and the relation of the latter to the former – to make sense. Incoherent combinations of attitudes, in a very real way, don't make sense, and nor does deliberation that takes them seriously. Respecting the demands of coherence isn't merely a matter of “making pretty patterns”;<sup>54</sup> such a description demeans the project of engaging in deliberation that makes sense. And incoherence is more than unpretty; it is (in a sense) unintelligible.

My view, then, is that the hegemony of value in the theory of normativity is a tyranny, and that we should resist it. But making out the broader case for this is a task for another day.

## Acknowledgments

For helpful discussions, questions, and comments, I'm grateful to Tom Dougherty, Zoë Johnson King, Benjamin Kiesewetter, Alex King, Stephanie Leary, Sarah Stroud, Adam Waggoner, Daniel Wodak, Michael Young, and audiences at Boston University, the Union College Workshop on Belief and Reasoning, the UNC-Tübingen Conference, and the University of St Andrews. I'm also grateful to participants in my graduate seminar at UNC in Fall 2019 for very helpful feedback. Special thanks to Chris Howard for numerous extremely helpful conversations that were formative in developing the ideas of the paper, and to Sarah Buss, Daniel Fogal, Wooram Lee, and two anonymous referees for *Noûs* for detailed written comments.

---

<sup>52</sup> Cf., e.g., Maguire (2016, 2018).

<sup>53</sup> Including Kolodny and Kiesewetter; see the references in fn. 24 above.

<sup>54</sup> Cf. Wedgwood (2017: 4).

## References

- Bratman, M. (2009). Intention, practical rationality, and self-governance. *Ethics*, 119, 411-443. doi: 10.1086/599985
- Broome, J. (1999). Normative requirements. *Ratio*, 12, 398-419. doi: 10.1111/1467-9329.00101
- Broome, J. (2004). Reasons. In R.J. Wallace, P. Pettit, S. Scheffler & M. Smith (Eds.), *Reason and Value: Themes from the Moral Philosophy of Joseph Raz* (pp. 28-55). Oxford: Oxford University Press.
- Broome, J. (2005a). Does rationality give us reasons? *Philosophical Issues*, 15, 321-337. doi: 10.1111/j.1533-6077.2005.00069.x
- Broome, J. (2005b). Have we reason to do as rationality requires? A comment on Raz. *Journal of Ethics and Social Philosophy*, 1(1) (Symposium), 1-10. doi: 10.26556/jesp.v1i1.13
- Broome, J. (2013). *Rationality Through Reasoning*. Chichester: Wiley-Blackwell.
- Brunero, J. (2012). Instrumental rationality, symmetry and scope. *Philosophical Studies*, 157, 125-140. doi: 10.1007/s11098-010-9622-0
- Chang, R. (2013). Grounding practical normativity: going hybrid. *Philosophical Studies*, 164, 163-187. doi: 10.1007/s11098-013-0092-z
- Dancy, J. (1977). The logical conscience. *Analysis*, 37, 81-84. doi: 10.1093/analys/37.2.81
- Dancy, J. (2000). Should we pass the buck? *Royal Institute of Philosophy Supplement*, 47, 159-173. doi: 10.1017/S1358246100006974
- Fogal, D. (2020). Rational requirements and the primacy of pressure. *Mind*, 129, 1033-1070. doi: 10.1093/mind/fzz038
- Fogal, D. & Worsnip, A. (forthcoming). Which reasons? Which rationality? *Ergo*.
- Greenspan, P. (1975). Conditional oughts and hypothetical imperatives. *Journal of Philosophy*, 72, 259-276. doi: 10.2307/2024734
- Heuer, U. (2011). Beyond wrong reasons: The buck-passing account of value. In M. Brady (Ed.), *New Waves in Metaethics* (pp. 166-184). Basingstoke: Palgrave-Macmillan.
- Hieronymi, P. (2013). The use of reasons in thought (and the use of earmarks in arguments). *Ethics*, 124, 114-127. doi: 10.1086/671402
- Howard, C. (2019). The fundamentality of fit. *Oxford Studies in Metaethics*, 14, 216-236. doi: 10.1093/oso/9780198841449.003.0010
- Johnson King, Z. (2019). We can have our buck and pass it too. *Oxford Studies in Metaethics*, 14, 167-188. doi: 10.1093/oso/9780198841449.003.0008
- Kiesewetter, B. (2017). *The Normativity of Rationality*. Oxford: Oxford University Press.
- Kolodny, N. (2005). Why be rational? *Mind*, 114, 509-563. doi: 10.1093/mind/fzi509
- Kolodny, N. (2007). How does coherence matter? *Proceedings of the Aristotelian Society*, 107, 229-263. doi: 10.1111/j.1467-9264.2007.00220.x
- Korsgaard, C. (1996). *The Sources of Normativity*. Cambridge: Cambridge University Press.
- Lord, E. (2018). *The Importance of Being Rational*. Oxford: Oxford University Press.
- Maguire, B. (2016). "The value-based theory of reasons," *Ergo*, 3, 233-262. doi: 10.3998/ergo.12405314.0003.009
- Maguire, B. (2018). There are no reasons for affective attitudes. *Mind*, 127, 779-805. doi: 10.1093/mind/fzx011

- Mildenberger, C. (2019). A reason to be rational. *Inquiry*, 9-10, 1008-1032. doi: 10.1080/0020174X.2018.1470570
- Neta, R. (2015). Coherence and deontology. *Philosophical Perspectives*, 29, 284-304. doi: 10.1111/phpe.12066
- Parfit, D. (2011). *On What Matters, Volume 1*. Oxford: Oxford University Press.
- Prichard, H.A. (1912). Does moral philosophy rest on a mistake? *Mind*, 21, 21-37. doi: 10.1093/mind/XXI.81.21
- Raz, J. (2005). "The myth of instrumental rationality," *Journal of Ethics and Social Philosophy*, 1/1, 2-28. doi: 10.26556/jesp.v1i1.1
- Rowland, R. (2014). Dissolving the wrong kind of reason problem. *Philosophical Studies*, 172, 1455-1474. doi: 10.1007/s11098-014-0359-z
- Scanlon, T. (2007). Structural irrationality. In G. Brennan, R. Goodin, F. Jackson & M. Smith (Eds.), *Common Minds: Themes from the Philosophy of Philip Pettit* (pp. 84-103). Oxford: Oxford University Press.
- Schroeder, M. (2004). The scope of instrumental reason. *Philosophical Perspectives*, 18: 337-364. doi: 10.1111/j.1520-8583.2004.00032.x
- Schroeder, M. (2007). The Humean theory of reasons. *Oxford Studies in Metaethics*, 2, 195-219.
- Schroeder, M. (2010). Value and the right kind of reason. *Oxford Studies in Metaethics*, 5, 25-55.
- Schroeder, M. (2012). The ubiquity of state-given reasons. *Ethics*, 122, 457-488. doi: 10.1086/664753
- Shah, N. (2006). A new argument for evidentialism. *The Philosophical Quarterly*, 56, 481-498. doi: 10.1111/j.1467-9213.2006.454.x
- Southwood, N. (2008). Vindicating the normativity of rationality. *Ethics*, 119, 9-30. doi: 10.1086/592586
- Wallace, R.J. (2014). Practical reason. *Stanford Encyclopedia of Philosophy*. <https://plato.stanford.edu/entries/practical-reason/>
- Way, J. (2010). Defending the wide-scope approach to instrumental reason. *Philosophical Studies*, 147, 213-233. doi: 10.1007/s11098-008-9277-2
- Way, J. (2011). The symmetry of rational requirements. *Philosophical Studies*, 155, 227-239. doi: 10.1007/s11098-010-9563-7
- Way, J. (2012). Transmission and the wrong kind of reason. *Ethics*, 122, 489-515. doi: 10.1086/664749
- Wedgwood, R. (2017). *The Value of Rationality*. Oxford: Oxford University Press.
- Worsnip, A. (2018a). The conflict of evidence and coherence," *Philosophy and Phenomenological Research*, 96, 3-44. doi: 10.1111/phpr.12246
- Worsnip, A. (2018b). Reasons, rationality, reasoning: how much pulling-apart? *Problema*, 12, 59-93.
- Worsnip, A. (forthcoming). *Fitting Things Together: Coherence and the Demands of Structural Rationality*. New York: Oxford University Press.
- Zimmerman, M.J. (2007). The good and the right. *Utilitas*, 19, 326-353. doi: 10.1017/S0953820807002622