

Rational Requirements and the Primacy of Pressure¹

Daniel Fogal
New York University

1. Two Dimensions of Rational Evaluation

I take there to be two threads in our thought and talk about rationality, both practical and theoretical. In one sense, rationality is roughly a matter of responding correctly to the reasons one has. This is the sense in which it's rational to believe that the sun will rise tomorrow or that grass is green, but irrational to believe that you're smarter than Einstein or that the moon is made of cheese, since the evidence had by nearly everyone—including you—strongly supports both of the former but neither of the latter. Likewise, it's rational to exercise regularly as well as look both ways before crossing the street, but irrational to spend your life memorizing telephone numbers or bet everything you have on a fair coin landing heads. In this sense, to be rational is to be reasonable. Call this *substantive rationality*.

In another sense, however, rationality is a matter of coherence, or having the right structural relations hold between one's attitudinal mental states—one's beliefs, intentions, preferences, hopes, fears, and the like—independently of whether those states (henceforth, 'attitudes') are reasonable or justified.² This is the sense in which it's rational to believe the obvious consequences of other things you believe and take the means you believe to be required to achieve your ends, while it's irrational to believe something you think is unsupported by the evidence or intend to do something you believe you shouldn't do. The relevant notion of coherence is thus a broad and broadly normative one, encompassing a range of different combinations of attitudes that intuitively clash, or fail to properly 'fit' together, where the lack of fit needn't involve logical inconsistency. Call this *structural rationality*.³

¹ Forthcoming in *Mind*. Thanks to Paul Boghossian, John Broome, John Brunero, Cian Dorr, Billy Dunaway, Matti Eklund, Jane Friedman, Matthew Frise, Martin Glazier, Camil Golub, Daniel Greco, Daniel Harris, Daniel Hoek, Jordan MacKenzie, Barry Maguire, Aidan McGlynn, Errol Lord, Andrew Reisner, Bernhard Salow, Sam Scheffler, Stephen Schiffer, Sam Shpall, Kurt Sylvan, Ang Tong, Dan Waxman, Jonathan Way, Roger White, Daniel Whiting, Alex Worsnip, Crispin Wright, and audiences at M.I.T., Rutgers University, Southampton University, University of Edinburgh, University of St. Andrews, Uppsala University, University of Bern, the St. Louis Annual Conference for Reasons and Rationality, several anonymous referees, the editors, and especially Jim Pryor for discussion and/or comments.

² My use of 'attitudes' is restricted to contentful mental states that are apt candidates for rational assessment, and thus excludes states, such as bodily sensations and perceptual experiences, that may play a justificatory role without themselves being assessable as rational or irrational, justified or unjustified, and so on.

³ Cf. Scanlon (2007), Worsnip (2018a,b,c), and Pryor (2018), among others.

Although the distinction between substantive and structural rationality has become increasingly commonplace, it remains controversial.⁴ It's a distinction that's motivated by intuitive judgments concerning particular cases together with the more general observation—emphasized by Kolodny (2007, 2008)—that there seems to be something *distinctively wrong* with incoherent sets of attitudes and something *distinctively right* with coherent ones, where this isn't simply a matter of such (sets of) attitudes being reasonable or unreasonable.⁵ I won't be arguing for it here, however.

The focus of this paper is on the nature of structural rationality. According to the now-standard view—one most prominently associated with John Broome (1999, 2007a,b, 2013)—structural rationality is associated with a distinctive set of *rules* or *requirements* that mandate or prohibit certain combinations of attitudes, and it's in virtue of violating these requirements that incoherent agents are structurally irrational. Call this the *requirements-based account* of structural rationality.⁷ Standard examples include consistency requirements (in imperatival form: don't believe contradictions! don't intend incompatible things!), instrumental requirements (intend the means you think are necessary to your ends!), closure requirements (believe the obvious consequences of other things you believe!), and enkratic requirements (intend to do what you believe you ought to do!). Structural requirements thus involve cognitive attitudes like belief and directive attitudes like intentions, as well as combinations of cognitive and directive attitudes. And just as there are requirements governing 'full' or 'outright' attitudes like belief and intention, so there are requirements governing 'partial' or 'graded' attitudes like credences and partial intentions.⁸ For simplicity, though, my focus will be on the former.⁹

I think the requirements-based account of structural rationality is mistaken. The goal of this paper is to explain why, and to motivate an alternative account. According to the view I favor, rather than corresponding to a distinctive set of law-like requirements, structural rationality should be seen as corresponding to a distinctive

⁴ Those who deny the distinction and/or its theoretical significance include Broome (2004, 2013), Kiesewetter (2017), Lord (2017, 2018), Parfit (2011), Schroeder (2008, 2011), and Wedgwood (2017), though the reasons why vary. I should note I do *not* intend the substantive/structural distinction to be an 'objective'/'subjective' or 'external'/'internal' distinction. One might take facts about substantive rationality *and* facts about structural rationality to obtain in virtue of facts about our non-factive mental states, for example, and hence be an internalist about both.

⁵ Although Kolodny himself wants to deny that such appearances are veridical (he grants there are enkratic requirements, but that's it), he recognizes the dialectical burden is more than just a matter of explaining why something has gone wrong in cases of incoherence.

⁷ For elaboration of the requirements-based view, as well as suggestions for how to develop it, see Fogal (2018b). Parts of this section and the next draw on the discussion there.

⁸ See Holton (2008) for the need to recognize partial intentions in addition to full ones.

⁹ I say 'requirements-based' rather than 'rules-based' since I'll be following the literature in focusing primarily on requirements. But requirements are only one type of rule, or principle; there may also be principles of permission (cf. Broome 2013, Ch. 10).

kind of pro tanto *rational pressure* or *force*—i.e. something that comes in degrees, having both magnitude (strength) and direction (for/against). Call this the *pressure-based account* of structural rationality. Something similar is standardly assumed to be true of substantive rationality. On the resulting picture, then, each dimension of rational evaluation is associated with a distinct (and distinctive) kind of rational pressure—substantive rationality with what I’ll call *justificatory* pressure and structural rationality with *attitudinal* pressure. The former is generated by one’s reasons while the latter is generated by one’s (mere) attitudes.

The pressure-based account of structural rationality is not completely new. Pryor (2004, 2018), for instance, suggests something similar in the theoretical realm. Unlike me, however, he doesn’t generalize it to the practical realm nor does he argue against the requirements-based view. The pressure-based alternative has otherwise been overlooked. As a result, the contemporary debate over the nature of structural rationality has proceeded on the largely unquestioned assumption that it bottoms out in requirements, rather than something else.

The same is not true of substantive rationality. In fact, it’s the reverse. That’s because we already standardly think of reason (mass noun) and justification—two of the normative notions at the heart of substantive rationality—as being force-like, having something akin to magnitude and direction. Reason and justification, after all, are both *gradable*—you can have more or less reason/justification to do something—and inherently *directional*—reason/justification is always reason/justification *to* do something or *not* to do something, and so always plays either a supporting or opposing role. (Here as elsewhere talk of ‘doing something’ is meant to include forming, maintaining, and revising attitudes, not just actions.) Reasons (count noun) are also inherently directional, and although they don’t themselves admit of degrees—something either is a reason or it isn’t—they are intimately related to something that does—namely, their normative significance or ‘weight’, with some reasons being weightier than others.¹⁰ Reason(s) and justification are also standardly assumed to be pro tanto, rather than strict, normative notions. That is, you’re not automatically irrational in virtue of not doing what you have reason to do, since it may be opposed or defeated, and you may have equal or more reason to do something else.¹¹

Structural requirements, in contrast, are *not* force-like—they don’t come in degrees—and are not standardly conceived of as being pro tanto.¹² Instead, they are strict, all-or-nothing affairs. They are strict in the sense that you are automatically irrational in virtue of violating them, and they are all-or-nothing in the sense that, for any action or attitude (or combination of such) *A*, either *A* is mandated or prohibited or it’s not. Although some rules or requirements may be more important than others, and hence take precedent in cases of conflict, to say that requirements admit of

¹⁰ For more on the subtleties and complexities of ‘reason(s)’-talk, see Fogal (2016, 2018a).

¹¹ Of course, whether we’re *correct* in think thinking of justification and reason(s) in force-like ways, and whether we’re correct in thinking of substantive rationality in terms of them, is a separate and important question—one I needn’t take a stand on.

¹² Brunero (2010) is a notable exception.

hierarchical relationships, such as rankings, is not to say that they admit of degrees. Rank-ability is not gradability.

Suffice it to say, when theorizing about a given domain it's a substantive question whether any particular bit of ideology is well-suited to the task. Ideology isn't innocent—it needs to be faithful to the phenomenon. Indeed, one of the most important trends in 20th century moral theory involves a change in ideology: the transition away from strict, threshold-y notions such as absolute principles (cf. Kant 1785/2012) to non-strict, threshold-y notions such as *prima facie* duties (cf. Ross 1930), and then to non-strict, graded notions like reason(s) (cf. Dancy 1993, 2000).¹³ When it comes to structural rationality, however, the ideology of requirements has rarely been called into question. By subjecting it to scrutiny, and suggesting a possible replacement, this paper serves as a corrective.

The plan for the remainder is as follows. In §2, I spell out the requirements-based account of structural rationality in more detail. In §3, I argue against it. The main objection is straightforward: all structural requirements appear susceptible to counterexample. Although I don't take the objection to be decisive, I do think it presents a serious challenge to the standard account. In §4, I sketch the alternative, pressure-based account, noting various similarities as well as differences between justificatory pressure (substantive rationality) and attitudinal pressure (structural rationality). Insofar as the pressure-based account can explain everything the requirements-based account can, and more besides, it's to be preferred.

2. The Requirements-Based Account of Structural Rationality

Central to the debate over structural rationality is the intuitive observation that the attitudes we actually have—whether or not they're justified—make a difference concerning what other attitudes it would be structurally rational, or coherent, for us to have. (Note that, for ease of expression, I'll often drop the 'structural(ly)' qualifier in what follows. Also, talk of 'having attitudes' is meant to include the absence as well as the presence of attitudes.) In particular, it seems that you can be *rationally committed* to having certain attitudes given that you have certain other attitudes, in such a way that you'll exhibit a rational failing if you have the latter without the former. If you intend to make dinner, for example, and believe that to do so you need to buy pasta, it would be irrational for you to not intend to buy pasta. Similarly, it seems you can be *rationally prohibited* from having certain attitudes given that you have certain other attitudes, in such a way that you'll exhibit a rational failing if you have the former while also having the latter. If you believe it's cold outside, for example, it would be irrational for you to believe it's *not* cold outside. For this reason, many philosophers take structural rationality to correspond, at least in part, to a distinctive set of *rules* or *requirements* that mandate or prohibit certain (combinations of) attitudes.

The most prominent and influential proponent of a requirements-based account of structural rationality is John Broome (1999, 2004, 2007a,b, 2013), but he's just one of

¹³ For more on this transition, and its importance, see Hurka (2014).

many who take structural requirements seriously. The basic idea is that for each kind of incoherent combination of attitudes there is a corresponding rule prohibiting it, and that what's wrong with incoherent agents is that they violate these rules. An analogy with the law is instructive. Suppose you fail to stop at a red light. What you did was illegal. Subtleties aside, the reason is clear: there's a law that requires you (and everyone else) to stop at red lights, and you failed to stop. If there hadn't been such a law, your failure to stop may have been reckless but it wouldn't have been illegal. In general, for each kind of illegal action there's a law prohibiting it, and it's in virtue of violating such laws that particular actions are illegal.

The requirements of structural rationality are supposed to play an analogous role: just as laws help explain why particular actions are legal or illegal, so structural requirements are supposed to help explain why particular (combinations of) attitudes are rational or irrational. Suppose someone—call him 'Tom'—believes that he's Superman, and that Superman can fly, but lacks confidence in his own ability to fly. Tom is structurally irrational. But why exactly? According to the requirements-based picture, it's because he fails to believe the obvious consequences of other things he believes, and there's a requirement prohibiting him (and us) from doing so. Structural requirements are thus supposed to be more than mere necessary conditions for being fully rational, and in at least two ways. First, the requirements, when violated, are supposed to guarantee a negative evaluative property—namely, being irrational—and not just the lack of a positive one—namely, being rational. Second, the requirements are supposed to state conditions such that agents who fail to meet them are irrational *in virtue of* doing so. Violating a requirement doesn't merely *guarantee* that you're irrational: it *explains why* you're irrational. The same goes for compliance: just as it is in virtue of violating such requirements that agents are irrational, it is in virtue of complying with them that they are rational.

Different authors express this point in different ways. For instance, Jonathan Way (2018) draws a distinction between stronger and weaker senses in which one might be 'rationally required' to do something:

In [a] weak sense, to say that you are rationally required to do A is to say that doing A is a necessary condition of being fully rational. However, [there's] a stronger sense in which [it might be thought that] rationality requires coherence. What, we might ask, explains why [deductive] incoherence and means-end incoherence are irrational? A natural answer is that there are *rules* or *principles* which require you to be closure and means-end coherent (...) If you have an incoherent combination of attitudes you are irrational *because* you violate a rational requirement. (Way 2018, p. 486-487)

In a similar vein, Broome (2007a,b, 2013) distinguishes between two senses of 'requires' and its cognates: a weaker 'property' sense and a stronger 'source' sense. He thinks the most interesting questions concern what rationality requires in the latter sense, not the former. Schroeder (2013) agrees, equating his own talk of the 'rules' of rationality with Broome's notion of a source requirement and noting that the debate concerns which rules you are violating when you are incoherent.

I'm focusing on this point about the explanatory status of the requirements of structural rationality because it's essential to the claim I want to reject.¹⁴ It would be foolhardy to deny there are necessary conditions of being fully rational, and in *that* sense requirements of rationality. After all, in order to be rational one has to have mental states, but having mental states isn't rationally required in the relevant sense.¹⁵ Instead, what I'm concerned to deny is that the existence of structural requirements in the *strong* sense. The nature of structural rationality—just like the nature of substantive rationality—is importantly *unlike* that of the law.

Despite the consensus concerning the status of structural requirements, there remains controversy over much else, including their exact content, their normative status, and whether they are 'wide' or 'narrow' scope. Although these issues are important, I won't be addressing them in what follows.¹⁶ That's because I wish to cast doubt on the logically prior claim that structural rationality is to be understood in terms requirements at all. The next section (§3) explains why.

3. Against Structural Requirements

My main objection to the requirements-based conception of structural rationality is simple: all structural requirements are susceptible to counterexample, and so false. To illustrate, I'll consider some examples from both the theoretical and practical realms.

3.1 All Structural Requirements are False

Let's begin with the following deductive coherence ('DC') principle, which is about as good a candidate as any for expressing a structural requirement applying to everyone, always, and of necessity:¹⁷

¹⁴ Kiesewetter (2017) and Lord (2017, 2018) are also skeptical of structural requirements, although their skepticism (unlike mine) is rooted in skepticism of structural rationality as such, together with optimism that their accounts of substantive rationality can explain everything that needs to be explained when it comes to our judgments of (ir)rationality. For reasons to be less optimistic, see Littlejohn (2018), Worsnip (2018c), and Fogal and Worsnip (MS).

¹⁵ Cf. Broome (2007b, 2013). Some philosophers, however, are only concerned with providing necessary conditions, and hence with what rationality requires in the weak sense—see, e.g., Titelbaum (2015) and Easwaran and Fitelson (2015). Nonetheless, for each non-trivial necessary condition, there remains the question of *why* it's necessary.

¹⁶ For discussion of these issues, see Broome (2013), Kiesewetter (2017), Lord (2018), and Brunero (MS). The normative status of structural requirements is a particularly important (and fraught) issue. Although I reject the existence of such requirements, a parallel question concerning the normative status of attitudinal pressure arises. I am unable to adequately address the issue here, however, and so leave it as an open question.

¹⁷ (DC) is formulated as a 'wide-scope' requirement rather than 'narrow-scope' one—it prohibits a combination of attitudes rather than a particular attitude. Nothing hinges on this—the narrow-scope version is similarly susceptible to counterexample.

(DC) Rationally requires that you not: believe p , believe if p then q , and fail to believe q .

So understood, I think (DC) is false. For suppose that although you believe p and that if p then q , you also happen to believe that you're really bad at making deductive inferences, or that you've been slipped a drug that severely impairs your reasoning ability.¹⁸ Alternatively, suppose that as a result of a spurious philosophical argument your confidence in the validity of modus ponens is shaken, and so you remain unsure about whether q despite believing p and that if p then q . In cases like these it's not clear what, if anything, you're required to believe, structurally speaking. Nonetheless, the presence of such higher-order attitudes makes a difference, and in a way that casts doubt on (DC). In particular, it doesn't seem like you're *definitely doing something wrong* in virtue of believing p , believing if p then q , and failing to believe q in such cases—contra (DC) understood as a strict, exception-less requirement in the 'strong' sense.

Similarly, suppose you happen to believe that you lack sufficient reason to believe p , or to believe that if p then q , or to believe q , or to believe any of them. Although it may be that you're irrational in virtue of having beliefs that you believe you shouldn't have—and hence being epistemically 'akratic'—it doesn't seem like you're definitely doing something wrong in virtue of believing p , believing if p then q , and failing to believe q . That misidentifies the locus of incoherence in such cases, which is between one's first-order beliefs and higher-order beliefs. Once again (DC) seems to overreach, since it does more than merely entail that one is irrational whenever it is violated—it also entails that one is irrational *in virtue of* violating it.

Alternatively, suppose that in addition to believing p and that if p then q , you also believe r and that if r then not- q . Or suppose you believe s and that if s then not- p . What then? Clearly you're incoherent, and not as you should be, structurally speaking. But although you've definitely done something wrong, you don't seem to have done something wrong in virtue of failing to believe q while believing p and believing if p then q . Here as before (DC) misidentifies the locus of incoherence, which is more global—it's the full set of conflicting attitudes, not any particular subset of them.

The defender of (DC) might respond by noting that this is a case where the relevant requirement applies, and is violated, twice. That is, both of the following combinations of attitudes violate the (schematic) requirement expressed by (DC):

{belief that p , belief that if p then q , failing to believe q }
{belief that r , belief that if r then not- q , failing to believe not- q }

The defender of (DC) can then point out that no matter which doxastic attitude you adopt vis-a-vis q , you're guaranteed to violate the requirement—assuming you retain

¹⁸ Cf. Christensen (2010), who offers analogous examples involving higher-order evidence, as opposed to merely higher-order attitudes.

your existing beliefs—and this fact is ‘global’ insofar as it concerns all the relevant attitudes.

Although correct as far as it goes, this response fails to capture the fact that there’s something *distinctively* bad about believing p and that if p then q while also believing r and that if r then not- q , over and above the fact that it guarantees that (DC) is violated twice over. That is, there’s something bad about *that very combination of attitudes*, not just one or more sub-combinations. So something like the following should be a requirement as well:

(DC+) Rationality requires that you not: believe p , believe if p then q , believe r , and believe if r then not- q .

But if that’s right, then you’ll be guilty of irrationality not just twice but *three* times over—twice in virtue of violating (DC) in two distinct ways and once in virtue of violating (DC+). And that seems wrong. You are definitely irrational in virtue of having an incoherent set of attitudes, but you don’t seem to be guilty of irrationality three times over. That’s over-counting.

The defender of (DC) might argue that the third violation is derivative from the first two, and so the agent has only gone wrong in two *basic* ways. But this fails to capture the sense that there’s something *distinctively* bad about the larger set of attitudes. Contrast, for instance, the following:

(DC++) Rationality requires that you not: believe p , believe if p then q , fail to believe q , and believe r .

Unlike (DC+), there’s nothing *distinctively* wrong with the combination of attitudes in (DC++). Although there’s *something* wrong with such a combination, the badness is traceable to the first three attitudes alone—the fourth (the belief that r) is extraneous. So unlike (DC+), the plausibility of (DC++) clearly derives from (DC).

To take another example from the theoretical realm, consider the following (‘NC’ stands for ‘no contradiction’):

(NC) You are rationally required to not: believe p and not- p .

As with (DC), there are grounds for rejecting (NC) as a strict, law-like requirement. Perhaps the most promising counterexample involves so-called ‘glut theorists’, who think there are true contradictions. As Jc Beall notes:

One notable sort of incoherence is often tied to logical inconsistency: rationality instructs us to reject (logical) contradictions—reject any sentence (or proposition, etc.) of the form $A \wedge \neg A$. But even this sort of principle needs to be balanced with the pursuit of increasing coherence. It may be, for example, that glut theorists are right: given conservativeness with respect to (say) truth principles or the like, the most coherent response to standard antinomies (e.g., liar paradox) takes them to be gluts. But such is the messy—and ‘defeasible’—life of rational inquiry. (Beall 2013, p. 3)

The point is not that glut theorists are right—I doubt they are. The point is rather that it seems at least *possible* for a sophisticated glut theorist to violate (NC) and yet not be structurally irrational, or guilty of incoherence, in virtue of doing so. It's one thing to be mistaken; it's another to be irrational.¹⁹

Call cases involving higher-order attitudes 'higher-order cases' and cases involving conflicting attitudes 'conflict cases'. My claim is that for any candidate requirement, there will be a range of either higher-order or conflict cases (or both) that constitute counterexamples to that requirement.²⁰ There's nothing particularly special about (DC) and (NC)—analogous considerations apply to other candidate requirements. The recipe for generating counterexamples is straightforward. For every level *n* or kind *K* of rationally evaluable attitude (belief, intention, preference, and so on), it's possible to adopt a range of higher-order attitudes concerning the unreliability and/or normative defectiveness of one's attitudes at level *n* or of kind *K*, and for any rationally evaluable attitude *A*, there are a range of other attitudes, or combinations of attitudes, that directly conflict with *A*. The main ingredient in coming up with counterexamples at each successive stage is simply creativity.

It's worth reiterating that I'm *not* denying there are necessary conditions for being fully or ideally coherent, and in *that* sense some requirements of structural rationality. Perhaps anyone who violates (DC) or (NC) is guaranteed to fall short of some rational ideal. Even so, that's not what the debate over structural requirements is about. Instead, the relevant requirements are supposed to be rules or principles which not only ensure that a subject is irrational whenever they're violated, but also that the subject is irrational *in virtue of* violating them. They're supposed to tell us something about the nature of structural rationality, not merely provide a diagnostic for it. *That's* the picture of structural rationality I'm resisting, and that the examples involving the downward impact of higher-order attitudes and the horizontal impact of conflicting attitudes are intended to cast doubt upon. If such cases reveal it's possible to violate (DC) or (NC) without doing *anything* wrong structurally, so much the better. (DC) and (NC) are two of the most plausible requirements, and so we might expect others to share the same fate—that is, we might expect that, for any candidate requirement (in the strong sense), there will be cases in which a subject violates it without being structurally irrational at all. But all I need is for higher-order and/or conflict cases to make plausible the weaker claim that for any given candidate requirement, it's possible to violate it without being irrational *in virtue of* violating it.

¹⁹ Broome (2013, p. 91) considers a similar example. He recognizes its intuitive force, and the challenge it poses. He then proceeds to consider an alternative, 'more liberal' account of the conditions of rationality that views rationality as 'meeting your own standards, so that you are not rational only if you fail by your own lights' (Broome 2013, p. 92). Broome rightly rejects such a view, but fails to consider other alternatives.

²⁰ Although I've been ignoring the possibility of there being principles of permission in addition to requirements, the same moral will apply to the former, at least insofar as they purport to be exception-less.

To see how the same dialectic arises with (purported) requirements of practical rationality, consider the following means-end (ME) principle:²¹

(ME) Rationality requires that you not: intend end *E*, believe *M* is a necessary means to *E*, and fail to intend *M*.

Understood as a strict requirement, I think (ME) is false. For suppose that although you intend *E* and believe *M* is a necessary means to *E*, you happen to believe that it's impossible for you to *M*, or that *M* would have terrible consequences for your friends and family. Or suppose you happen to believe that you shouldn't intend *E*, or that the evidence doesn't support believing that *M* is necessary for *E*. In cases like these, although you're clearly incoherent, and not as you should be, structurally speaking, it doesn't seem like you're incoherent in virtue of failing to intend *M* while intending *E* and believing *M* is necessary for *E*.

Alternatively, suppose that in addition to intending *E* and believing *M* is a necessary means to *E*, you also intend *E2* and believe that not-*M* is a necessary means to *E2*. Or suppose you intend *E3* and believe *E3* is incompatible with *E*. What then? Although you have clearly gone astray, structurally speaking, you don't seem to have gone astray in failing to intend *M* while intending *E* and believing *M* is necessary for *E*. Here as before (ME) misidentifies the locus of incoherence, which is more global. Of course, as with (DC), the defender of (ME) might grant that there is a more global source of incoherence while nonetheless insisting that (ME) identifies an additional, self-standing source. But again as with (DC), this gives rise to over-counting worries.

For another example from the practical realm, consider the following principle prohibiting 'strongly' inconsistent intentions (SII):

(SII) Rationality requires that you not: intend *A*, intend *B*, and believe that it is impossible to both *A* and *B*.

Bratman (1984) appeals to a principle like (SII) in objecting to the 'Simple View' of intentional action, according to which intentionally acting requires an intention to so act. He asks us to imagine someone—call him Bruce—playing two video games simultaneously, one with each hand. The goal of each game is to guide a missile and hit a target, and there is a reward for hitting either one. However, the games are connected in a way that makes it impossible to hit both—the machines will shut down if both targets are about to be hit. Bruce knows this, but continues to play both games in order to increase his chances of hitting a target. Now suppose that Bruce manages to hit target 1 as a result of his effort and skill. It seems he hits the target intentionally (cf. Bratman 1984, p. 382). But did he *intend* to hit target 1? Not if (SII) is true. For if Bruce intended to hit target 1, he must have also intended to hit target 2—after all, he

²¹ As with most requirements, there's controversy concerning its details. Here and elsewhere, I'm ignoring possible complications in the interest of simplicity—none of the details I'm aware of are sufficient to thwart the threat of counterexample.

was ‘trying equally hard, and with equal skill, as well as with equally weak confidence of success, to hit target 2’ (Bratman 1984, p. 382). Since he knows it’s impossible to hit both, he violates (SII). So if (SII) is true, it follows that Bruce is structurally irrational. But intuitively he’s not. So something has to give—either Bruce did not intend to hit target 1 or else (SII) is false. Bratman opts for the former—on his view, Bruce hits the target intentionally without intending to hit the target.²² But I find it more plausible to reject (SII), and to view the video game case as a case in which recognizably inconsistent intentions are structurally rational, just as the so-called preface paradox plausibly illustrates the possibility of recognizably inconsistent yet structurally rational beliefs.²³

As a final illustration, consider the following anti-akrasia (AA) requirement:

(AA) Rationality requires that you not: believe that you ought to *A* and fail to intend *A*.

Here again higher-order and conflict cases provide grounds for rejecting (AA). Brunero (2013) makes a version of this point in arguing for the possibility of rational akrasia. He asks us to consider a person, Jack, who believes that he ought to punish his daughter even though this belief conflicts with ‘many other beliefs and desires he has’, including the belief that he lacks sufficient evidence for the claim that he ought to punish his daughter:²⁴

[Unfortunately,] Jack knows that he’s unable to change this belief. He knows that even if he were to review all the convincing arguments against punishing her, and remind himself that there’s not sufficient evidence for his belief, he would continue to believe that he ought to punish his daughter. (...) If, in light of this knowledge, Jack decides not to punish his daughter, Jack would have the combination of attitudes prohibited by [(AA)], but yet Jack’s having this combination wouldn’t be [structurally] irrational. (Brunero 2013, p. 553-554)

Brunero is careful not to claim that Jack is fully rational, structurally speaking, since ‘a fully rational person would not hold beliefs he believes to be insufficiently supported by the evidence’ (Brunero 2013, p. 554). But as he rightly points out:

[I]t doesn’t follow from Jack’s not being fully rational that Jack’s *akrasia* is irrational. For one thing, we can locate the specific source of Jack’s irrationality and it’s not his

²² For fuller responses to Bratman, see McCann (1991) and Sverdlik (1996).

²³ Bratman considers but rejects what I take to be the correct view, one vindicated by the pressure-based view—namely, that while there is a ‘general presumption’ against strong inconsistency, this presumption is ‘overridden in the present case’ (Bratman 1984, p. 386).

²⁴ Brunero’s example is a modified version of one offered by Audi (1990). Unlike Audi, Brunero is clear that he’s concerned with structural, as opposed to substantive, rationality (neither uses those terms).

akrasia, but his *theoretical* irrationality. For another, Jack’s not intending to punish his daughter seems to be a rational way of *responding* to this irrational, recalcitrant belief. Even though a fully rational version of Jack wouldn’t violate [(AA)], it seems wrong to accuse *Jack* of irrationality in violating [(AA)]. (Brunero 2013, p. 554)

More generally, we shouldn’t be surprised at the possibility of a gap arising between ‘what one’s fully [structurally] rational self would do’ and ‘what it would be [structurally] rational for one’s actual self to do, given the ways in which one’s actual self falls short of full rationality’ (Brunero 2013, p. 555).

3.2 *Ceteris Paribus to the Rescue?*

I’ve now considered and rejected five central principles of structural rationality, construed as strict requirements in the strong sense. As far as I can tell, a similar dialectic arises in relation to all other candidates, thereby casting doubt on the requirements-based conception of structural rationality in general. I don’t pretend to have settled the matter, however. The discussion so far merely represents an opening salvo—one that is highly suggestive, but not conclusive.

For one thing, it might be objected that the above dialectic merely reveals the need to add further conditions to (DC), (ME), and their kin—such as the absence of higher-order funny business, the absence of fancy philosophical views, and the absence of other conflicting attitudes. Even so, without something like a *ceteris paribus* clause, the continual threat of counterexample will likely result in structural requirements of considerable (perhaps never-ending) complexity, with the result looking increasingly ad hoc and less explanatory at each step. Rather than adding conditions to handle each potential counterexample, then, the defender of a requirement like (DC) might try to add a ‘no defeaters’ or ‘*ceteris paribus*’ clause. Consider, for instance:

- | | |
|--------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------|
| (DC) _{ND} | Rationality requires you, in the absence of disruptive attitudes, to not: believe <i>p</i> , believe if <i>p</i> then <i>q</i> , and not believe <i>q</i> . |
| (DC) _{CP} | <i>Ceteris paribus</i> , rationality requires you to not: believe <i>p</i> , believe if <i>p</i> then <i>q</i> , and not believe <i>q</i> . |

Although plausible as general principles, (DC)_{ND} and (DC)_{CP} are less plausible when construed as explanatory, law-like requirements. After all, much of the explanatorily and normatively important action will occur ‘off-stage’ in determining whether the principles apply at all. With (DC)_{ND} it’s in determining whether there are any normatively disruptive attitudes, and with (DC)_{CP} it’s in determining whether other things are equal (whatever that amounts to). Such principles, when they apply, simply assume the relevant normative action has been resolved one way rather than another, without themselves saying how or why. Even when they apply, then, such principles presuppose underlying normative facts that in turn require explanation.

This then raises the specter of explanatory irrelevance: for whatever explains why the principles apply, when they do, would also seem to be capable—together with the

agent's antecedent beliefs—of explaining why the agent shouldn't believe p , believe if p then q , and not believe q . That is, it seems reasonable to expect a 'common cause' explanation of the general principle together with the particular cases it concerns, with whatever it is that explains why principles like $(DC)_{ND}$ and $(DC)_{ND}$ apply, when they do, also being what ultimately explains why agents shouldn't believe p , believe if p then q , and not believe q , when they shouldn't. In particular, it would seem to be facts concerning the totality of the agent's attitudes that would explain both the absence of normatively disruptive action as well the irrationality of having such a combination of attitudes under such circumstances. $(DC)_{ND}$ and $(DC)_{ND}$ themselves would thus seem to be at best explanatorily derivative and at worst epiphenomenal.

Indeed, this is just what we should expect on the pressure-based account of structural rationality I favor. On that view, explanations of structural rationality, like those of substantive rationality, are fundamentally 'bottom up' rather than 'top down': when an agent is irrational in virtue of being incoherent, it is ultimately in virtue of failing to respond correctly to the attitudes they have—or, rather, to the balance of the attitudinal pressure generated by their attitudes—just as when an agent is irrational in virtue of being unreasonable, it is ultimately in virtue of failing to respond correctly to the reasons they have—or, rather, to the balance of the justificatory pressure generated by their reasons. Principles like $(DC)_{ND}$ or $(DC)_{ND}$ are thus viewed as true generalizations that are made true by explanatorily more fundamental facts—namely, facts about the existence of and interaction between various attitudinal (as opposed to justificatory) pressures. They summarize facts about structural rationality rather than explain them. On the pressure-based view, the same is true of principles like (DC) and (ME) . The next section (§4) explains why.

4. The Pressure-Based Account of Structural Rationality

4.1 *Motivating the Pressure-Based Account*

There are at least two reasons why the pressure-based account of structural rationality should be taken seriously. The first is that it avoids the difficulties facing the requirements-based view outlined above (§3)—a point I'll return to below. The second is that it enjoys intuitive support. Indeed, if anything, it deserves to be treated as the default. For as Kolodny observes in motivating the debate over 'requirements of formal coherence as such':

The intuitive idea is that formally incoherent attitudes give rise to a certain *normative tension*, or exert a kind of *rational pressure* on one another, and this tension, or pressure, is relieved, just when one of the attitudes is revised. (Kolodny 2008, p. 437, emphasis added)

Although Kolodny makes similarly suggestive remarks elsewhere—including in his seminal ‘Why be rational?’ (2005) paper—he nonetheless never seems to recognize how uneasily they fit with a requirements-based view.²⁵ Or consider Pryor:

[T]he mere belief that your color vision is defective—whether justified or not—would rationally obstruct you from believing the wall is red on the basis of your color experiences. It would *rationaly pressure* you to place less trust in those grounds on that question. (Pryor 2004, p. 364, emphasis added)

Similar claims are found scattered throughout the literature on structural rationality, though Pryor is the only one I’m aware of who is sensitive not only to the naturalness but also to the distinctiveness of thinking about structural rationality in pro tanto, gradable (i.e. pressure-like) terms.

At a minimum, quotes like these testify to the fact that appeals to requirements (and other law-like principles) fail to exhaust our intuitions about structural rationality. This is reinforced by reflection on examples. Suppose, for instance, that a juror has heard from the prosecution but not the defense. Although she believes (perhaps unreasonably) that there’s considerable evidence indicating that the defendant is guilty, she’s not sure how much evidence the defense will present in support of the defendant’s innocence. In such a case there seems to be at least *some* rational pressure for the juror to believe the defendant is guilty. But it would be premature to have reached a settled opinion on the matter, and so it’s not the case that, structurally speaking, she ought or is required to believe that the defendant is guilty. This is even clearer if the juror happens to believe (again, perhaps unreasonably) that the local police and the prosecutor are thoroughly corrupt, and hence that the defendant might have been framed. While there remains some pressure to believe the defendant is guilty, there’s also pressure to believe the opposite. Given the conflicting pressures, the appropriate response is to suspend judgment.

Analogous cases can be constructed in the practical realm concerning the relationship between what you believe it would be good to do and what it would be structurally rational for you to (intend to) do. Consider, for example, Huckleberry Finn, who befriends Jim (a slave) and helps him escape. As Arpaly (2002) explains,

Huckleberry is plagued by what he calls ‘conscience.’ He believes (...) that helping a slave escape amounts to stealing, and stealing is wrong. He also believes that one should be helpful and loyal to one’s friends, but loyalty to friends is outweighed by some things, such as property rights, and does Miss Watson, Jim’s owner, not have property rights? (...) Yet when the opportunity comes to turn Jim in, [Huckleberry] experiences a strong reluctance to do so, [largely due to the] fact that he has come to perceive Jim as a person. (Arpaly 2002, pp. 228, 229-230)

²⁵ The same goes for Brunero (2010). See also Ferrero (2016).

Clearly Huckleberry is in a conflicted state of mind. Given his belief that it is wrong to steal and that Jim ought to be returned to Miss Watson, there seems to be rational pressure for him to intend to turn Jim in. But given his belief that one should be helpful and loyal to one's friends as well as his tacit (but nonetheless real) belief that Jim is a "full person, just like Huckleberry himself" (Arpaly 2002, p. 229), there seems to be rational pressure for him to *not* intend to turn Jim in.²⁶ Such pressures conflict, and neither seems decisive on its own.

More generally, claims like the following seem plausible:²⁷

| | |
|-----------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Evidence | If you believe there's (at least some) evidence in favor of p , then there's (at least some) pro tanto rational pressure to believe p . |
| Goodness | If you believe that it would be good (in some way, to some degree) for you to ϕ , then there's (at least some) pro tanto rational pressure to intend to ϕ . |

Neither Evidence nor Goodness specifies *how much* pressure there is in a given case. Plausibly, however, it's partly a function of how good you take the evidence or action to be (independently of how good it *actually is*), as well as how confident you are in that judgment. It also depends on the impact of other possible sources of attitudinal pressure – Evidence and Goodness specify one such source (namely, beliefs about evidence and beliefs about goodness), but there are many others. The amount or strength of pressure generated by a particular attitude or combination of attitudes therefore varies, and won't always (or even usually) be strong enough to *require* anything of you all by itself.²⁸ The fact that the stringency of the normative relations

²⁶ Arpaly is primarily concerned with the moral status of Huckleberry's subsequent refusal to turn Jim in; I'm primarily concerned with its rational status.

²⁷ There may be exceptions. Whether there are, and if so why, will depend (inter alia) on substantive issues concerning the different kinds of attitudes we have, how they're related, and what their rational import is, exactly. For example, we might think of resolutions or, more generally, all-things-considered practical judgments (whether justified or not) as having a special kind of rational force that effectively silences any potentially competing attitudinal pressures generated by, say, our beliefs about what it would be good (in some way, to some degree) to do. Such complexities are compatible with a pressure-based account of structural rationality. My goal in this paper, however, is merely to provide a broad outline of such an account, leaving a range of important intramural issues intentionally unresolved.

²⁸ This is compatible with the relevant pressures requiring something of you when there are no competing pressures, or when competing pressures are undermined, outweighed, or otherwise defeated. As I emphasize later, it is the *overall balance* of attitudinal pressure—not any particular amount or source of such—that determines what you structurally ought or are required to do.

associated with structural rationality comes in degrees suggests that judgments about what is structurally required in particular cases capture only part of the story.

Perhaps surprisingly, the most prominent defender of a requirement-based account, John Broome, has expressed sympathy with something like this point in the past. In his early and seminal work on structural requirements, Broome (1999) introduces a distinct operator to represent what is *recommended* but not necessarily *required*, on analogy with the distinction between merely ‘pro tanto’ reasons and ‘decisive’ reasons. Yet reference to recommendation-like principles all but disappears from his subsequent work, which focuses primarily on requirements.²⁹ Even so, merely introducing a second operator that is ‘slack’ rather than ‘strict’ is expressively inadequate in the case of structural rationality—it doesn’t allow us to draw fine-grained enough distinctions. Introducing a new operator corresponding to each degree of ‘slackness’ would not only lead to an unsightly proliferation of operators, but it would also obscure the fundamental continuity of the underlying phenomenon and fail to reflect the various complex interactions of attitudinal support, opposition, and undermining I highlight below in §4.2.

In what remains I’ll present the alternative, pressure-based account of structural rationality in more depth. I’ll begin by comparing it to the standard, pressure-based account of substantive rationality (§4.2). I’ll then suggest a way in which substantive and structural rationality can be seen as unified (§4.3) despite being ultimately distinct (§4.4). I’ll then conclude by explaining how one can explain the appeal of general structural principles without taking them to be strict requirements (§4.5).

4.2 Similarities Between Attitudinal and Justificatory Pressure

As noted at the outset, rather than being seen as corresponding to a distinctive set of *requirements* or *rules*, I take each dimension of rational evaluation to correspond to a distinct kind of pro tanto rational *pressure* or *force*—substantive rationality with *justificatory pressure* and structural rationality with *attitudinal pressure*. On this view, believing p and that if p then q generates substantial attitudinal pressure to also believe q , just as having good reason to believe p and that if p then q generates (or, rather, constitutes) substantial justificatory pressure to believe q . Similarly, intending to φ and believing that φ -ing requires ψ -ing generates substantial attitudinal pressure to intend to ψ , just as having good reason to φ and good reason to believe φ -ing requires ψ -ing generates (or constitutes) substantial justificatory pressure to intend to ψ . Analogous claims hold for partial or graded attitudes.

²⁹ Primarily but not exclusively: Broome (2013) also introduces a special class of ‘basing’ principles, some of which prohibit certain transitions between attitudes (he calls these ‘basing prohibitions’) and others of which permit certain transitions (‘basing permissions’). These are diachronic principles connecting an agent’s attitudes at one time with their attitudes at another time, and so not of present concern. But even so, basing prohibitions and permissions are both essentially ‘threshold-y’ or ‘all-or-nothing’, and hence importantly unlike ‘graded’ or ‘quantitative’ normative notions such as value, justification, reason, and (on my view) attitudinal pressure, all of which come in degrees.

Of course, we're used to thinking of one's reason(s)/justification/evidence as either being or being closely associated with something force-like—something with *magnitude* (strength), which can vary, and *direction* (for/against). That much is old news. What I'm suggesting is that we recognize a *distinct* kind of pro tanto rational pressure—one that behaves similarly but is importantly unlike one's reason(s)/justification/evidence insofar as it is generated by the *mere presence* of our beliefs, intentions, and the like, regardless of their justificatory standing. It too has something like magnitude (strength), which can vary, and direction (for/against). And it exhibits the same kind of complex intra-dimensional interactions that justificatory pressure does, all of which come in degrees: just as there might be some justificatory/attitudinal pressure that *supports* an attitude or action, there might also be some justificatory/attitudinal pressure that *opposes* it, and still other justificatory/attitudinal pressure that plays an *undermining* role, weakening—or, at the limit, eliminating—other justificatory/attitudinal pressures.

To illustrate the possibility of undermining: just as acquiring evidence that a friend is a habitual liar undermines (to some degree) whatever justification you may have had to believe what he says, so merely *believing* that he's a habitual liar reduces (to some degree) whatever pre-existing attitudinal pressure there may have been to believe what he says. To nonetheless continue trusting your friend without hesitation would be irrational, though the nature of the irrationality in the two scenarios is different. Hence the sense in which both pressures are not just pro tanto (or capable of being 'opposed') but also underminable (capable of being 'undermined')—or, in a word, *defeasible* (capable of being 'defeated' in either way). The force metaphor, with the concomitant distinction between component forces and resultant force, allows for different ways of understanding the nature of such interactions. In cases involving opposition, for instance, we can think of the initial component pressure as remaining constant even though the overall resultant pressure is altered by competing component pressures, whereas in cases involving undermining we can think of the initial component pressure as itself being weakened (and at the limit eliminated), and of the resultant pressure changing as a result. So even if the resultant pressure ends up being the same in both cases, the explanation of why each ends up that way differs.

The force metaphor highlights another important similarity between substantive and structural rationality: just as it's plausible to take facts about what you're justified in doing (believing, intending, etc.) to be determined by facts about justificatory pressure, rather than vice versa, so it's plausible to take facts about what you're structurally required and/or permitted to do to be determined by facts about attitudinal pressure, rather than vice versa. On the view I favor, then, facts about rational pressure are taken to be normatively and explanatorily fundamental with respect to both substantive *and* structural rationality, and facts about normative statuses like being required or adequately justified are taken to be resultant, derivative facts. To say this is to make an explanatory claim, of a broadly metaphysical sort: in both cases the relevant 'threshold-y' normative facts obtain in virtue of the corresponding 'contributory' facts, not vice versa.

This allows us to say something more specific about the property of being fully rational. For as Broome notes, '[p]erhaps the most important question a system of rational requirements needs to settle is whether you are rational—have the property of rationality' (Broome 2007a, p. 363). Requirements-based theorists like Broome have a straightforward answer: you are (fully) rational if and only if you satisfy all the requirements of rationality that apply to you. Since I reject the requirements-based account of rationality, I don't face the question Broome considers. Nonetheless, an analogous question arises within the pressure-based framework concerning the relationship between rational pressure and the person-level property of being rational. By 'rational' Broome means what I mean by 'structurally rational'. But the issue arises with substantive rationality as well. So Broome's question bifurcates—there are two kinds of rational pressure (justificatory and attitudinal) and two corresponding properties (being substantively rational and being structurally rational), and we need to say something about the relationship between each.

Strictly speaking, the pressure-based account is compatible with a variety of answers. But here's what I take to be the most natural account: just as being substantively rational is a matter of doing (believing, intending, etc.) what the overall balance of justificatory pressure supports—or, more colloquially, what you have most reason to do—and doing so in the right way, so being structurally rational is a matter of doing what the balance of attitudinal pressure supports, and doing so in the right way. In neither case is an attitude's being rational a matter of there being more or less pressure to adopt it; nor is it a matter of an agent responding correctly to any individual amount or source of pressure considered on its own. Instead, what matters is being properly responsive to how the various pressures balance out. In both cases, the person-level property of being rational is to be understood in terms of all-things-considered facts about rational pressure, together with a proper basing condition. These global, all-things-considered facts about rational pressure are in turn determined by the complex interactions between various local, graded facts about (pro tanto) rational pressure.

It's worth emphasizing that although attitudinal pressure is generated by or grounded in facts about our psychology, it's not itself psychological.³⁰ In general, we shouldn't confuse normative facts and relations with what grounds them. What's more, I take the source of the pressure to arise from one's having the attitude itself—and *not* from one's thinking there are good reasons for having it, or from any other attitude about that attitude. The presence of such higher-order attitudes may well make a difference, however, since they will be sources of attitudinal pressure.

³⁰ There may nevertheless be a close connection between them—cf. Worsnip (2018b). Also, I should note that I'm not assuming every being capable of having propositional attitudes is thereby subject to attitudinal pressure. Although whether one has propositional attitudes—and if so, which kinds—is important, whether and how attitudinal pressure is generated may also depend on one's cognitive capacities (cf. Fink 2014). As Broome (2013) and others emphasize, only suitably complex beings are apt targets for rational evaluation. Just how complex they need to be, and in what ways, is a hard question.

Suppose, for instance, you believe p and if p then q , but come to think q isn't supported by the evidence. In such a case the former attitudes, taken together, generate pressure to believe q , while the latter generates pressure to *not* believe q , and hence together with the belief that if p then q generates pressure to not believe p . How the competing attitudinal pressures are resolved will depend, in part, on their respective strengths.

It's also worth emphasizing that because attitudinal pressure is always grounded in the presence of various attitudes, the *failure* to have a certain attitude is not itself a source of attitudinal pressure. (Here I'm departing from the earlier convention of using 'attitude' to include absences of attitudes.) So whereas believing p and if p then q generates attitudinal pressure to also believe q , believing if p then q while failing to believe q doesn't generate attitudinal pressure to not believe p . As a result, if you believe p and believe if p then q , but haven't seriously considered the question of whether q (you haven't put 'two and two together' yet), it follows that there's attitudinal pressure to believe q , but it *doesn't* follow that there's any attitudinal pressure to not believe p . It's only if one were to, say, suspend judgment with respect to q , or view q as unlikely, or otherwise doubt that q that such attitudinal pressure would arise.³¹

4.3 *The Underlying Unity of Attitudinal and Justificatory Pressure*

So far I've emphasized two (related) points of similarity between justificatory and attitudinal pressure: both are force-like, having something akin to magnitude and directionality, and both enter into a range of complex relations, such as support, opposition, and undermining. Another important similarity is encapsulated by what I'll call the 'proportionality thesis'. According to this thesis, the strength and direction of attitudinal pressure generated by the presence of a given attitude (or combination of attitudes) will, at least in general, be proportional to whatever the strength and direction of justificatory pressure there would be were one to be fully justified in having that attitude (or combination of attitudes), while holding everything else fixed. Recall Tom, for example, who believes that he's Superman and that Superman can fly. How much attitudinal pressure is there for him to believe that he himself can fly, given his existing beliefs? Well, consider: how much justificatory pressure would there be for him to believe that he can fly, were he to be fully justified in believing that he's Superman and that Superman can fly, holding everything else fixed? The answer seems to be the same in both cases—the pressure, though defeasible, is quite strong. Though perfect precision shouldn't be expected, the indeterminacy in our judgments would lessen were we to find out how confident Tom is in each proposition. The relation between belief and confidence is of course complicated, but it's generally agreed that one can believe p and believe q while nonetheless being more confident in one than the other. In general, the more confident one is the more attitudinal pressure there will be, just as the more justification one has the more justificatory pressure

³¹ For a defense of suspended judgment as a positive, irreducible doxastic attitude, see Friedman (2013a,b).

there will be. Attitudinal pressure can thus be understood—at least heuristically—as a kind of ‘hypothetical’ justificatory pressure.³²

However, this gives rise to a natural worry concerning the purportedly sharp distinction between kinds of rational pressure. For given the systematic parallels between justificatory and attitudinal pressure, one might reasonably doubt they are entirely distinct things, and instead hope to understand one in terms of the other, or both in terms of some third thing. The most obvious way to explain the similarities, after all, would be to show that they are ultimately a unified phenomenon.

I can meet the objector halfway. For although I take substantive and structural rationality to be distinct and autonomous—in ways and for reasons to be explained shortly (§4.4)—it’s nonetheless plausible to view substantive rationality as enjoying a certain kind of explanatory priority. The fact that the amount and direction of attitudinal pressure generated by particular (combinations of) attitudes seems to track facts about the amount and direction of justificatory pressure there would be were the subject to be justified in having those attitudes (holding everything else fixed) suggests that structural rationality behaves as something like the ‘shadow’ of substantive rationality. It’s as if the primary rational imperative is to have justified attitudes (be reasonable!), but given that we’re imperfectly rational creatures—all of us have at least some ‘bad’ attitudes—there’s a subsidiary imperative to at least have a set of attitudes that could, in principle, be justified. Even when our attitudes are not *actually* justified, then, rationality still expects us to function cognitively *as if* they’re justified.³⁴ So there may well be a sense in which substantive rationality enjoys explanatory priority, and in which there’s an underlying unity between the two dimensions of rational evaluation. What I’m skeptical of, however, is any attempt to deny the existence of either or to fully reduce one to the other.³⁵

4.4 The Distinctness of Attitudinal and Justificatory Pressure

Although attitudinal and justificatory pressure display important similarities, they are nonetheless distinct. They are distinct at least in the sense that they have different sources—attitudinal pressure is generated by (mere) attitudes while justificatory pressure is not (at least not in general). Arguably, however, they are also distinct in

³² Cf. Pryor (2018), who opts for talk of ‘hypothetical’ normative support and offers (what I call) the proportionality thesis as a heuristic in the theoretical realm. So-called ‘subjective reasons’ views (e.g. Parfit 1997, 2011; Schroeder 2008, 2011; Whiting 2014) are similar in structure, but nonetheless importantly different in a number of ways.

³⁴ This provides at least a partial answer to the question of the normative status of structural rationality. See Williamson (forthcoming) and Lasonen-Aarnio (forthcoming) for recent epistemological discussion of the distinction between primary and derivative norms, and Sylvan (2018) for discussion of derivative norms of respect.

³⁵ This leaves open a range of important metanormative questions concerning structural rationality (cf. Fink 2014, Worsnip 2018b). Many of them, however, are ones that advocates of the pressure-based view can disagree on, and so not ones that I’m directly concerned with in this paper.

the sense that their verdicts can come apart. One can imagine a die-hard conspiracy theorist, for example, who forms and maintains a complicated yet coherent web of beliefs in part by being utterly dismissive of any potentially conflicting evidence or testimony. The degree to which he is structurally rational far exceeds the degree to which he is substantively rational.

The possibility of being substantively rational (to a certain degree) without being structurally rational (to that degree) is also plausible, though more controversial. One prominent illustration involves permissive cases—that is, cases where it’s permissible by the lights of substantive rationality to adopt one of a number of different attitudes and yet structurally impermissible to adopt them all.³⁶ For example:

[S]uppose that one is at a restaurant and that there are three available dishes, each delicious in their own, different way. It’s plausible that in some such cases, one’s reasons are permissive with respect to one’s preferences: no particular pairwise preference would be substantively irrational. From that, it follows that there’s no substantive irrationality in preferring dish A to dish B, or in preferring dish B to dish C, or in preferring dish C to dish A. But it’s structurally irrational to have all of these preferences together. [I]t seems that there is structural irrationality without any guarantee of substantive irrationality. (Worsnip 2018c; cf. Quinn 1990)

Whether analogous cases arise in the theoretical domain depends on whether, given one’s evidence, there is ever more than one doxastic attitude it would be permissible to adopt.³⁸ If, as seems plausible, there are at least *some* cases in which your evidence gives you at least *some* leeway concerning which doxastic attitude to adopt (say, credence .55 or .56 in *p*), then we can construct cases in which an agent has a set of justified but incoherent doxastic attitudes (say, credence .55 in *p* and .44 in *not-p*). The large literatures surrounding the ‘preface’ and ‘lottery’ paradoxes also contain examples in which it seems reasonable for a subject to adopt a set of recognizably inconsistent doxastic attitudes that are at least *prima facie* structurally irrational.³⁹

Another illustration of the possibility of substantive rationality without structural rationality involves misleading higher-order evidence.⁴⁰ Suppose, for example, I have decisive reason to finish an overdue paper, though a trusted advisor with an excellent track record of advice recommends taking a vacation instead. It seems at least possible for her testimony to carry enough evidential weight to justify me in believing I should go on vacation. If I were to then adopt the attitudes it would be most substantively rational for me to adopt, I would end up intending to do something I

³⁶ Cf. Way (2013, 2018), who notes that the same point can be made with cases involving incommensurability and supererogation.

³⁸ See Kelly (2014) and White (2014) for a recent exchange.

³⁹ For examples, see Christensen (2004) and Easwaran and Fitelson (2015), among others.

⁴⁰ Cf. Worsnip (2018a).

believed I shouldn't do, and so be akratic—a paradigm of structural irrationality (though defeasibly so).⁴¹ Similar cases can be constructed in the theoretical realm.⁴²

Although there's controversy concerning each of the cases I've considered, taken together they make plausible the possibility of having attitudes that are substantively rational (to a certain degree) without being structurally rational (to that degree), and thus the possibility of substantive and structural rationality coming apart in *both* directions. In addition to being distinct, however, I take substantive and structural rationality to be *autonomous* in the sense that the presence of one kind of pressure doesn't—by itself—guarantee the presence of the other kind of pressure. Although there might be circumstances that guarantee the presence of both kinds of pressure, on the assumption that *mere* attitudes do not generally generate justification—unlike, say, experiences, which do—the explanation of why one kind of pressure exists will be distinct from that of the other. I also take substantive rationality and structural rationality to be autonomous in the further sense of failing to directly interact with each other. For instance, if there's substantial justificatory pressure to believe *p* (e.g. you have lots of evidence supporting *p*), this pressure isn't in any way reduced simply because you happen to believe something obviously incompatible with *p*. Your mere *belief* doesn't, by itself, make any difference to what your evidence supports. Although I don't have a neat-and-tidy argument for such claims, I take them to be plausible enough to adopt as working hypotheses, and hence assume that there are interactions between pressures of the same kind without there being interactions between pressures of different kinds.⁴³

There are other possible ways in which substantive and structural rationality differ, in addition to being distinct and having distinct grounds. For one thing, attitudinal pressure is *escapable* in a way that justificatory pressure isn't, since one can revise one's attitudes but not one's reason(s)—at least not in the same way or to the same extent. That is, we have a certain kind of direct (though largely non-voluntary) control over the facts that generate attitudinal pressure—namely, our attitudes—that we lack over the grounds of justificatory pressure—namely, our reason(s). So even if you happen to believe *p* and if *p* then *q*, you can always in principle (if not in practice) make it *not* the case that you ought to believe *q* by revising one of your antecedent attitudes, since in doing so you remove the source of attitudinal pressure to believe *q*. The normative force of reason(s), however, is not so easily escaped. One might engage in further inquiry by doing more research, for instance, or consulting those who are in a better epistemic position, but these are typically extended, complex activities that involve interacting with others and other sources of information.⁴⁴

⁴¹ This is similar to the Arpaly case involving Huck Finn, though in discussing the latter I wasn't assuming that Huck's moral beliefs were epistemically justified.

⁴² Cf. Pryor (2004, 2018), Greco (2014), Horowitz (2014), and Worsnip (2018a).

⁴³ Pryor (2018), in contrast, thinks 'hypothetical and categorical norms [roughly: attitudinal and justificatory pressure] can combine to generate more hypothetical norms' (128).

⁴⁴ In this respect the view on offer is similar to that of Sam Shpall (2013, 2014), who offers an account of pro tanto rational 'commitments' (as opposed to reasons or requirements) on

Another potential difference between substantive and structural rationality concerns the determinacy and specificity of their all-things-considered verdicts. For while it is often assumed there will be a fairly determinate verdict concerning whether—and, if so, to what (rough) degree—some particular action or attitude is supported by the reason(s) one has, and hence whether one is justified in performing or having it, there won't always be a similar verdict in cases involving mere attitudes. For in cases of attitudinal conflict, such as ones involving inconsistent beliefs or intentions, the attitudinal pressures generated by the conflicting attitudes won't combine as happily as they do in cases of attitudinal concord, where there are no conflicting attitudes. In cases of attitudinal concord the attitudinal pressures will combine in much the same way as the justificatory pressures would combine were the attitudes to be justified (per the proportionality thesis), whereas in cases of attitudinal conflict the attitudinal pressures will *not* usually combine in such a way. That's because in cases of attitudinal conflict it'll normally be impossible for the conflicting attitudes to be jointly justified, and hence impossible for you to have a reasonable twin—a person who has all the same attitudes but, unlike you, is justified in having them. As a result, in cases of attitudinal conflict there normally won't be a fact of the matter concerning how the justificatory pressures would combine were those attitudes to be justified. And so even if there's a global 'all-things-considered' verdict concerning which particular attitude(s) the combined *justificatory* pressures support—and there might not, if (e.g.) reasons can have imprecise weights—we shouldn't assume there will also be one concerning what the combined *attitudinal* pressures support. Instead, we'll sometimes (perhaps often) have to rest content with local, pro tanto verdicts concerning the strength and direction of attitudinal pressure generated by particular (combinations of) attitudes, with the only global, all-things-considered verdicts being irreducibly disjunctive.”

A pair of examples will help illustrate the point. Recall Tom, who believes that he's Superman and that Superman can fly. Suppose first that despite believing that he's Superman and that Superman can fly, Tom hasn't seriously considered the question of whether he himself can fly. Upon considering it, however, he proceeds without hesitation to make the obvious inference. In this case, Tom's pre-existing beliefs generate substantial attitudinal pressure to believe that he can fly and so, given the absence of countervailing pressure, the most rational response, structurally speaking, is to believe he can fly. That's the response that's called for by his standing beliefs, and there's a fairly clear sense in which Tom believes as he should—even though there's another sense in which Tom does *not* believe as he should, since doing so runs contrary to all of his evidence.

which they're grounded in an agent's attitudes and hence similarly escapable. However, the same sort of considerations that cast doubt on the existence of rational requirements in the strong sense can be marshaled against the existence of Shpallian commitments. Neither one is needed to do important explanatory work, and their plausibility can be explained in terms of attitudinal pressure.

But now consider a variant of the example in which the possibility of flying strikes Tom as wildly implausible, whether justifiedly or not. And suppose his confidence in his inability to fly is comparable to his confidence in his identity as Superman and in Superman’s ability to fly. In this case, unlike the previous case, there doesn’t seem to be a clear verdict concerning what the most structurally rational response is overall. And that’s because the relevant pressures don’t combine to single out any particular attitude for revision. Any two of Tom’s beliefs taken together generate substantial attitudinal pressure to revise the third—hence the sense in which, given any two of his beliefs, he should revise the third—but no particular attitude is singled out for revision by the balance of attitudinal pressure itself. The tension or conflict can only be resolved by revising one of the beliefs in question. The most we can say is that, all things considered, Tom is structurally required to *not* have the three beliefs together.⁴⁵

Call the former ‘Coherent Tom’ and the latter ‘Incoherent Tom’. Taken together, they illustrate the point that although attitudinal pressures sometimes combine in such a way as to require (or otherwise support) particular attitudes—as with Coherent Tom—other times they combine in such a way as prohibit (or otherwise oppose) a certain combination of attitudes—as with Incoherent Tom. In general, on the pressure-based picture, cases of attitudinal concord will be cases in which a specific response is called for, while cases of attitudinal conflict will be ones in which something irreducibly disjunctive (i.e. a certain combination of attitudes) is structurally required. The directedness of attitudinal pressure and the variability in its verdicts is something that’s obscured by the focus on requirements.

4.5 Rational Pressure and Rational Requirements

It should be clear that although I’m skeptical of the explanatory import of principles like (DC) and (ME), I’m *not* skeptical about there being facts about what is required in particular circumstances, understood as all-pressure(-of-a-certain-type)-considered verdicts. For even if there are some plausible counterexample-free principles, it’s a further question as to whether they are genuinely explanatory in a ‘top-down’ fashion (as requirements in the strong sense are supposed to be) as opposed to merely being true generalizations—ones that are underwritten by facts about the rational pressures generated by the underlying facts (i.e. reasons and attitudes).

To illustrate, let’s return to (DC). How might we explain its plausibility by appealing to facts about attitudinal pressure? To begin with, consider the following pressure-based principles, neither of which purport to be strict, law-like requirements:

- (DC_p) If you believe *p* and believe if *p* then *q*, then there is substantial (pro tanto) attitudinal pressure to believe *q*.

⁴⁵ This is consistent with *justificatory pressures* singling out a particular attitude for revision; but that’s neither here nor there.

(ND_B) If there is substantial (pro tanto) attitudinal pressure to believe *p*, and no defeating pressure, then one is structurally required to believe *p*.

(DC_p) is true in virtue of the fact that believing *p* and believing if *p* then *q* generates substantial attitudinal pressure to believe *q*, and (ND_B) is a consequence of the general principle that you are structurally required to *A* just in case (and because) the overall balance of attitudinal pressure supports *A*. Notice that when evaluating (DC), it's natural to focus on the attitudes specified in the principle itself—i.e. the belief that *p*, the belief that if *p* then *q*, and the absence of belief in *q*—and fail to consider the possibility of conflicting and/or higher-order beliefs. Given the pressure-based picture, it should come as no surprise that such a combination of attitudes is judged to be irrational. After all, the beliefs specified in (DC) guarantee substantial attitudinal pressure to believe that *q* (cf. (DC_p)) and the absence of conflicting and higher-order beliefs guarantees that there won't be any countervailing, or otherwise defeating, pressure. Failing to believe *q* will therefore be structurally irrational (cf. (ND_B)). More generally, 'ordinary' cases in which you believe *p* and believe if *p* then *q* will be cases in which there is substantial undefeated attitudinal pressure to believe *q*, and hence cases in which failing to believe *q* is structurally irrational. Ordinarily, then, to avoid irrationality one must either (a) believe that *q*, (b) not believe that *p*, or (c) not believe that if *p* then *q*—the very combination of attitudes demanded by (DC). It's only in cases of higher-order funny business and fancy philosophical views that such a combination of attitudes has even the remotest chance of being structurally permissible. The pressure-based account is thus able to explain the plausibility—and, indeed, truth—of (DC) as a general, possibly-exception-permitting principle of structural rationality, despite its implausibility as a strict, law-like requirement.

A similar explanation is available of (ME). As with (DC), when evaluating (ME), it's natural to focus on the attitudes specified in the principle itself—namely, intending end *E*, believing *M* is a necessary means to *E*, and failing to intend *M*—and ignore the possibility of conflicting and/or higher-order attitudes. On the pressure-based picture, it should come as no surprise that such a combination of attitudes is judged to be irrational. Consider:

(ME_p) If you intend end *E* and believe *M* is a necessary means to *E*, then there is substantial (pro tanto) attitudinal pressure to intend *M*.

(ND_I) If there is substantial (pro tanto) attitudinal pressure to intend *A*, and no defeating pressure, then one is structurally required to intend *A*.

As (ME_p) makes clear, the attitudes specified in (ME) guarantee substantial attitudinal pressure to intend *M*, and the absence of conflicting and higher-order beliefs guarantees that there won't be any defeating pressure. Failing to intend *M* will therefore be structurally irrational, per (ND_I). More generally, 'ordinary' cases in which you intend *E* and believe *M* is a necessary means to *E* will be cases in which there is substantial undefeated attitudinal pressure to intend *M*, and hence cases in

which not intending M is structurally irrational. Ordinarily, then, to avoid irrationality one must either (a) intend M , (b) not intend E , or (c) not believe that M is a necessary means to E —the very combination of attitudes demanded by (ME). Hence the plausibility of (ME) as a general principle of structural rationality.

There's also a straightforward explanation of why ordinary conditionals of the form 'If you believe p and if p then q , then you should also believe q ' and 'If you intend to E and believe that in order to E you have to M , then you should intend to M ' seem true. It's because they *are* true, at least on one natural reading. On this reading, they're to be understood as claims about what is (normally) best, structurally speaking, given the truth of the antecedent.⁴⁶ Holding fixed the fact that you believe p and if p then q , for example, the structurally best response—in general, though not necessarily always—is to believe q . The other salient options, after all, are to either believe not- q or else suspend judgment with respect to q , both of which would normally result in incoherence. Similarly, holding fixed the fact that you intend E and believe that M is a necessary means to E , the structurally best response—in general, though again not always—is to intend M . The counterexamples to (DC) and (ME) when construed as strict law-like requirements thus fail to threaten the truth of their ordinary language counterparts—the counterexamples concern abnormal cases; ordinary conditionals ordinarily do not.

Explanations like those above generalize. The reason why, for example, believing propositions of the form p and not- p is normally structurally irrational is that doing so requires believing both p and believing not- p , and each of these manifestly incompatible beliefs generates substantial attitudinal pressure to revise the other.⁴⁷ It takes a lot of attitudinal fanciness to reduce or otherwise disrupt this tension, and in the absence of such disruption it will be structurally irrational to have both beliefs. Hence the plausibility of (NC). A similar explanation is available of (SII), with Bratman's video game case being exactly the kind of case involving attitudinal fanciness that is capable of rendering otherwise impermissible attitudes permissible.

The ability of the pressure-based view to provide such explanations serves to underscore the point that the existence of plausible general principles—as well as plausible natural language conditionals taken to express or support such principles—isn't enough to motivate positing requirements in the strong sense. Assuming the pressure-based view can explain the motivating data, a further argument in favor of positing law-like requirements is needed—one that defenders of the requirements-based view have yet to provide. Absent such an argument, while one can admit there are structural requirements in the *weak* sense, one should be wary of recognizing any in the strong sense. Instead, we can treat purported 'requirements' of structural

⁴⁶ For a fleshed-out version of this explanation, see Worsnip (2015).

⁴⁷ There may be cases in which their incompatibility is *not* manifest, if (for example) sentences like 'Superman can fly' and 'Clark Kent can fly' express the same proposition. In that case, 'Superman can fly but Clark Kent can't' would express a contradiction, one it might be rational to believe if one fails to realize Superman is Clark Kent. I wish to remain neutral, however, on the metaphysics of propositions, so I'm ignoring such complications.

rationality as true generalizations underwritten by more fundamental facts concerning rational pressure (of the attitudinal variety).

This is especially attractive if some of the relevant generalizations are exception-permitting. Whereas exceptions pose problems on the requirements-based account, they are to be expected on the pressure-based one. For in ‘higher-order cases’ the higher-order attitudes are just what they seem to be—namely, sources of attitudinal pressure that interacts with the various pressures generated by the lower-order attitudes. Likewise, in ‘conflict cases’ the relevant attitudes generate horizontal attitudinal pressures that conflict with each other. When you believe p and if p then q , for instance, you are thereby structurally pressured to believe q , and in the normal case—i.e. absent higher-order or conflicting attitudes—this pressure will be enough to require you to believe q . But in abnormal cases, the pressure may be discounted (as in higher-order cases), weakened (as in paradigmatic cases of undermining), or counteracted (as in cases of opposition). All of this can happen in various ways and to various degrees.

The emerging account of substantive and structural rationality is pleasingly economical: facts about rationality are explained in terms of two flavors (justificatory and attitudinal) of one kind of thing (pressure) rather than two flavors of two kinds of things (pressure and requirements). This allows us to avoid explanatory redundancies, as well as the question of how, if at all, pressures and requirements are supposed to interact or relate in particular cases.

5. Conclusion: The Primacy of Pressure

To sum up: on the view I’ve suggested, *pressure is primary*—that is, with respect to *both* dimensions of rational evaluation, facts about rational pressure are normatively and explanatorily prior to facts about what one rationally ought or is required to do. This stands in marked contrast to the dominant view of structural rationality, which takes requirements (and law-like principles more broadly) to occupy center stage, as well as to the tendency to take facts about what we structurally ought, or are required, to do as explanatorily basic.

Taking pressure to be primary opens up the possibility of rejecting the existence of any strict, ground-floor rational requirements, other than the general constitutive requirements to be reasonable and to be coherent (or at least not be incoherent). That’s all rationality in either sense directly and unconditionally requires of us at all times. In imperatival form, we’re required to

Respect justificatory pressure! (Be reasonable!)

Respect attitudinal pressure! (Be coherent!)

and that’s it. What we’re specifically required to do will vary, depending on the particularities of the situation we happen to be in. On this view, requirements themselves are at best a footnote in the theory of rationality.

References

- Arpaly, Nomy 2002, 'Moral Worth', *The Journal of Philosophy*, 99(5): 223-245.
- Audi, Robert 1990, 'Weakness of Will and Rational Action', *Australasian Journal of Philosophy*, 68(3): 270-281.
- Beall, Jc 2009, *Spandrels of Truth* (Oxford: Oxford University Press).
- Broome, John 1999, 'Normative Requirements', *Ratio*, 12(4): 398-419.
- 2004, 'Reasons', in R. J. Wallace, S. Scheffler, and M. Smith (eds), *Reason and Value: Essays on the Moral Philosophy of Joseph Raz* (Oxford: Oxford University Press).
- 2007a, 'Wide or Narrow Scope?', *Mind*, 116(462): 359-370.
- 2007b, 'Requirements', In *Hommage à Wlodek. Philosophical Papers Dedicated to Wlodek Rabinowicz*. www.fil.lu.se/hommageawlodek
- 2013, *Rationality through Reasoning* (Wiley-Blackwell).
- Brunero, John 2010, 'The Scope of Rational Requirements', *The Philosophical Quarterly*, 60(238): 28-49.
- 2013, 'Rational Akrasia', *Organon F*, 20(4): 546-566.
- MS, *Instrumental Rationality: The Normativity of Means-Ends Coherence* (under contract with Oxford University Press).
- Christensen, David 2004, *Putting Logic in its Place* (Oxford: Oxford University Press).
- 2010, 'Higher-Order Evidence', *Philosophy and Phenomenological Research*, 81(1): 185-215.
- Dancy, Jonathan 1993, *Moral Reasons* (Wiley-Blackwell).
- 2000, *Practical Reality* (Oxford: Oxford University Press).
- Easwaran, Kenny, and Branden Fitelson 2015, 'Accuracy, Coherence, and Evidence', in T. Szabó Gendler and J. Hawthorne (eds.), *Oxford Studies in Epistemology*, Vol. 5 (Oxford: Oxford University Press).
- Ferrero, Luca 2016, 'Pro-Tempore Disjunctive Intentions', in R. Altschuler and M. Sigrist (eds.), *Time and The Philosophy of Action* (Routledge).
- Fink, Julian 2014, 'A Constitutive Account of "Rationality Requires"', *Erkenntnis*, 79: 909-941.
- Fogal, Daniel 2016, 'Reasons, Reason, and Context', in E. Lord and B. Maguire (eds.), *Weighing Reasons* (Oxford: Oxford University Press).
- 2018a, 'Deflationary Pluralism about Motivating Reasons', in V. Mitova (ed.), *The Factive Turn in Epistemology* (Cambridge: Cambridge University Press).
- 2018b, 'On the Scope, Jurisdiction, and Application of Rationality and the Law', *Problema: Anuario de Filosofía y Teoría del Derecho*, Vol. 12: 21-57.
- Fogal, Daniel and Alex Worsnip MS, 'Which Reasons? Which Rationality?'
- Friedman, Jane 2013a, 'Suspended Judgment', *Philosophical Studies*, 162(2): 165-181.
- 2013b, 'Rational Agnosticism and Degrees of Belief', in T. Szabó Gendler and J. Hawthorne (eds.), *Oxford Studies in Epistemology*, Vol. 4 (Oxford: Oxford University Press).
- Greco, Daniel 2014, 'A Puzzle about Epistemic Akrasia', *Philosophical Studies*, 167(2): 201-19.
- Holton, Richard 2008, 'Partial Belief, Partial Intention', *Mind*, 117(465): 27-58.
- Horowitz, Sophie 2014, 'Epistemic Akrasia', *Noûs*, 48(4): 718-44.

- Hurka, Thomas 2014, *British Ethical Theorists from Sidgwick to Ewing* (Oxford: Oxford University Press).
- Kant, Immanuel 1785/2012, *Groundwork of the Metaphysics of Morals*, trans. Mary Gregor & Jens Timmermann (Cambridge: Cambridge University Press).
- Kelly, Thomas 2014, 'Evidence Can Be Permissive', in M. Steup, J. Turri, and E. Sosa (eds.), *Contemporary Debates in Epistemology*, 2nd ed. (Wiley-Blackwell).
- Kieseewetter, Benjamin 2017, *The Normativity of Rationality* (Oxford: Oxford University Press).
- Kolodny, Niko 2005, 'Why Be Rational?', *Mind*, 114(455): 509–563.
- 2007, 'How Does Coherence Matter?', *Proceedings of the Aristotelian Society*, 107: 229–263
- 2008, 'The Myth of Practical Consistency', *European Journal of Philosophy*, 16(3): 366–402.
- Lasonen-Aarnio, Maria forthcoming, 'Virtuous Failure and Victims of Deceit', in J. Dutant and F. Dorsch (eds.), *The New Evil Demon: New Essays on Knowledge, Justification, and Rationality* (Oxford: Oxford University Press).
- Littlejohn, Clayton 2018, 'Being More Realistic About Reasons: On Rationality and Reasons Perspectivism', *Philosophy and Phenomenological Research* (online first).
- Lord, Errol 2017, 'What You're Rationally Required to Do and What You Ought to Do (Are the Same Thing!)', *Mind*, 126(504): 1109–1154.
- 2018, *The Importance of Being Rational* (Oxford: Oxford University Press).
- McCann, Hugh 1991, 'Settled Objectives and Rational Constraints', *American Philosophical Quarterly*, 28(1): 25-36.
- Parfit, Derek 1997, 'Reasons and Motivation', *Proceedings of the Aristotelian Society, Supplementary Volumes*, 71: 99–146.
- 2011, *On What Matters*, Vol. I (Oxford: Oxford University Press).
- Pryor, James 2004, 'What's Wrong with Moore's Argument?' *Philosophical Issues*, 14(1): 349–378.
- 2018, 'The Merits of Incoherence', *Analytic Philosophy*, 50(1): 112-141.
- Quinn, Warren S. 1990, 'The Puzzle of the Self-Torturer', *Philosophical Studies*, 59(1): 79–90.
- Ross, W. D. 1930, *The Right and the Good* (Oxford: Oxford University Press).
- Scanlon, Thomas 2007, 'Structural Irrationality', in G. Brennan, R. Gooden, F. Jackson, and M. Smith (eds.), *Common Minds: Themes from the Philosophy of Philip Pettit* (Oxford: Oxford University Press).
- Schroeder, Mark 2008, 'Having Reasons', *Philosophical Studies*, 139(1): 57–71.
- 2011, 'What Does It Take to 'Have' a Reason?' in A. Reisner and A. Steglich-Peterson (eds.), *Reasons for Belief* (Cambridge: Cambridge University Press).
- 2013, 'Scope for Rational Autonomy', *Philosophical Issues*, 23(1): 297–310.
- 2014, 'Hypothetical Imperatives: Scope and Jurisdiction', in *Explaining the Reasons We Share* (Oxford: Oxford University Press).
- Shpall, Sam 2013, 'Wide and Narrow Scope', *Philosophical Studies*, 163(3): 717–36.
- 2014, 'Moral and Rational Commitment', *Philosophy and Phenomenological Research*, 88(1): 146–72.
- Sverdlik, Stephen 1996, 'Consistency Among Intentions and the 'Simple View'', *Canadian Journal of Philosophy*, 26(4): 515-522.
- Sylvan, Kurt 2018, 'Veritism Unswamped', *Mind*, 127(506): 381-435.

- Titelbaum, Michael 2015, 'Rationality's Fixed Point (or: In Defense of Right Reason)', in T. Szabó Gendler and J. Hawthorne (eds.), *Oxford Studies in Epistemology*, Vol. 5.
- Way, Jonathan 2013, 'Intentions, Akrasia, and Mere Permissibility', *Organon F*, 20(4): 588-611.
- 2018, 'Reasons and Rationality', in D. Star (ed.), *Oxford Handbook of Reasons and Normativity* (Oxford: Oxford University Press).
- White, Roger 2014, 'Evidence Cannot Be Permissive', in M. Steup, J. Turri, and E. Sosa (eds.), *Contemporary Debates in Epistemology*, 2nd Ed. (Wiley-Blackwell).
- Whiting, Daniel 2014, 'Keep Things in Perspective: Reasons, Rationality, and the A Priori', *Journal of Ethics and Social Philosophy*, 8: 1-22.
- Williamson, Timothy forthcoming, 'Justifications, excuses, and skeptical scenarios', in J. Dutant and F. Dorsch (eds.), *The New Evil Demon: New Essays on Knowledge, Justification, and Rationality* (Oxford: Oxford University Press).
- Worsnip, Alex 2015, 'Narrow-scoping for Wide-scopers', *Synthese*, 192(8): 2617–2646.
- 2018a, 'The Conflict of Evidence and Coherence', *Philosophy and Phenomenological Research*, 96(1): 3-44.
- 2018b, 'What is (In)coherence?', in R. Shafer-Landau (ed.) *Oxford Studies in Metaethics* Vol. 13 (Oxford: Oxford University Press).
- 2018c, 'Review of Keisewetter's *The Normativity of Rationality*. *Notre Dame Philosophical Reviews*'.