**[Note:** This is still very much a draft; please forgive the typos, incomplete references, missing material, etc. The introduction and bibliography are in a particularly sorry state, but will be fixed up soon. Any and all feedback is most welcome!**]**

## TABLE OF CONTENTS

# INTRODUCTION: The Importance of Normative Ideology

Quine (1951) famously distinguished between ontology and ideology:

> "Given a theory, one philosophically interesting aspect of it into which we can inquire is its ontology: what entities are the variables of quantification to range over if the theory is to hold true? Another no less important aspect into which we can inquire is its ideology (this seems the inevitable word, despite unwanted connotations): what ideas can be expressed in it?" (24)

One might object to Quine's narrow conception of ontology—there's plausibly more to ontology than just what exists[1]—but it's hard to deny that he was right to distinguish ontology from ideology and insist on the importance of the latter. A theory's ideology is its conceptual toolkit. At a minimum, it's a matter of what terms or concepts it employs in accounting for or explaining the relevant phenomena—i.e. that which it is a theory *of*—as well as what those terms or concepts mean—what their theoretical significance is. Ontology and ideology are distinct, but related. For an important criteria of a theory's adequacy is how well it captures the phenomena—how accurately and thoroughly it represents or provides an account of its subject matter. If a theory's ideology is radically impoverished vis-a-vis its subject matter—if the phenomenon is more complex, or different in kind, than any conceptual item the theory's toolkit—then the theory will inevitably fall short in its task of describing and explaining.

Of course, pretty much every theory falls short in some way. Our grasp of reality is partial and imperfect. But it's one thing for a theory to be *incomplete*; it's another for a theory to be *misguided*. Ideology isn't innocent. It needs to do justice to the phenomenon.

---

[1] Cf. Fine (2001), Schaffer (2009), and others.

Ideology changes. Sometimes a bit of ideology is to be *jettisoned* because it turns out not to be tracking anything in reality (phlogiston, witches), sometimes it needs to be *revised* or *replaced* because what it was supposed to be tracking is importantly different than originally thought (mass, simultaneity), and other times it is to be retained but *demoted* in status because although it tracks something in reality, what it tracks can be understood in other, more fundamental terms. Indeed, as Quine notes, as a subdivision of ideology there is "the question of what ideas are fundamental or primitive for a theory, and what ones derivative" (14).

Just as ideology matters in science and everyday life, so it matters in normative inquiry—in ethics, epistemology, and the rest. [Here I'll distinguish between strict vs. non-strict normative notions (e.g. all-things-considered-ought vs. pro tanto ought), threshold-y vs. graded normative notions (e.g. being justified vs. having some amount of justification), and others. I'll illustrate the distinctions and the different roles they played historically by focusing on one particularly prominent/importnat example from moral theory: absolute principles (Kant) vs. prima facie duties (Ross) vs. reasons (Dancy).]

The ideology appropriate to a particular normative domain may vary from case to case. And that's because there may be important differences in the nature and structure of different normative domains. Indeed, I strongly suspect there are, and that the differences break down in a principled way. In particular, I'm inclined to think that the domain of (what I'll stipulatively call) "natural normativity"—which includes the likes of morality, rationality, value, and prudence—invariably bottoms out in "contributory" or "weighted" normative facts—e.g. facts about what's good-making/bad-making, right-making/wrong-making, and reason-giving—that combine in certain complex ways to determine "overall" or "resultant" normative facts—e.g. facts about what one ought or is required to do or believe. In contrast, I think it's plausible that the domain of "artificial normativity"—which includes the likes of the

law, etiquette, and fashion —invariably bottoms out in "threshold-y" or "all-or-nothing" normative facts —e.g. facts about what is required or permitted, proper or improper, cool and uncool, and the like. Such a sorting is obviously highly controversial, but I think it is supported by the best account of the nature and structure of each individual sub-domain (morality, rationality, the law, etiquette, etc.). But that's not something I'll be concerned to argue for in this dissertation.[2] Instead, my goal will be to make incremental progress by focusing on the specific case of rationality and arguing against the dominant requirements-based—or, more generally, principles-based—account of structural rationality.

This dissertation will focus on two bits of ideology with considerable contemporary currency: the notion of a requirement of rationality and the notion of a normative reason. What I'll be arguing is that we should jettison the notion of a strict rational requirement and replace it (for now) with the more graded notion of (what I call) attitudinal pressure. And I'll argue that we should also be wary of the ideology of reasons, at least for the purpose of substantive normative (and metanormative) theorizing. In both cases the tendency to work with only one normative notion—one that is strict in the former case and insufficiently analyzed in the latter case—has obscured what really matters, both normatively and metaphysically.

---

[2] A related issue is why certain (sub-)domains of normativity, or standards of appraisal, seem to inherently matter in a way that others do not. In particular, it's often thought that the standards associated with natural normativity—such as morality, rationality, and prudence—are "robustly" and "inherently" normative in a way that the standards associated with artificial normativity—such as etiquette, fashion, and the law—are not. The question of what explains this contrast has been neglected, with metanormative theorists focusing primarily on providing accounts of natural (and "robust") normativity and giving artificial normativity little more than a passing glance (cf. Wodak). (The law is something of an exception.)

# CHAPTER 1: TWO DIMENSIONS OF RATIONALITY

## 1.1 A puzzle: conflicting oughts

Suppose you meet someone on the street who claims to be Superman. Suppose further that this person is perfectly sincere—he does, in fact, believe that he's Superman. It should be obvious that something is wrong with this person. Among other things, he has a crazy belief—he believes something that flies in the face of all the evidence. But suppose you also find out that despite believing that Superman can fly ("It's one of his greatest powers", he says) he is not at all sure that he himself can fly ("I gave up trying after my third broken leg"). Once again it should be obvious that something is wrong with him. Not only does he have a crazy belief, but he's also deductively incoherent in failing to believe the obvious consequences of other things he believes.

However, the second failing is interestingly different than the first, as evidenced by the seemingly paradoxical way we're prone to describe what's going wrong with the subject—call him 'Tom'.[3] On the one hand, it seems right to say that one should believe the obvious consequences of other things one believes, and so there's a sense in which Tom should believe that he can fly. On the other hand, one shouldn't believe something in the face of overwhelming evidence to the contrary, and so there's also a sense in which Tom should *not* believe that he can fly. We're thus faced with the puzzle of wanting to say both of these things—namely, that Tom should believe he can fly, and that he should not.

[3] This way of setting up the contrast is indebted to Setiya (2007), though the Superman example originated with Jim Pryor.

Similar examples involving strict means-end incoherence arise in the practical realm. Setiya (2007: 650) offers the following story (inspired by Rawls (1971)) to illustrate the "problem of instrumental reason":

> Imagine that I embark upon on a thoroughly irrational project: I intend to count the blades of grass in my garden…Despite my intention, however, I do not take what I know to be the necessary means. Even though I see that I have no chance to complete the enumeration unless I keep track of how many blades of grass I counted [and] where I counted them, I can't be bothered with bookkeeping. So, every morning, I am forced to start again [and never] complete the count.

As before, there are at least two things going wrong with such a subject—call her 'Jane'. On the one hand, given her goal it seems clear that Jane should be keeping track of the blades of grass. But on the other hand, Jane shouldn't be counting blades of grass in the first place, and so shouldn't be keeping track of them. Again we're faced with the puzzle of wanting to say both of these things—namely, that Jane should take the means that are necessary to achieve her end, and that she shouldn't.

Of course, apparently conflicting 'should'-judgments aren't always puzzling. The demands of morality, for instance, regularly conflict with the demands of self-interest, and there's nothing especially mysterious about the clash. What's interesting about cases like Tom and Jane is that the 'should'-judgments seem to arise from the same domain—they're both naturally understood as claims about what the *rational response* is in a given situation. To not believe the obvious consequences of other things you believe seems to constitute a *rational* failing; so does believing something in the face of overwhelming evidence to the contrary. The same is true of failing to intend the necessary means to one's end, and intending to do something you have no good reason to do.

It should be clear that we've hit upon a pattern, and that the foregoing observations generalize beyond the cases of deductive and means-end incoherence, both of which are extreme examples of an

otherwise pervasive phenomenon—one arising whenever there is a conflict or lack of "fit" between one's mental states or attitudes. (Throughout this dissertation I'll understand 'attitudes' broadly so as to include *absences* of attitudes as well.) This includes not just beliefs and intentions, but also hopes, fears, concerns, suppositions, worries, preferences, regrets, action-guiding desires, and the like.[4] I might realize it's more important to get a good night's sleep than to stay up late and catch up on the news, and yet prefer to continue reading. I might realize smoking is bad for me, yet not care about the consequences. I might know that spiders are mostly harmless, and yet still be afraid of them. It's possible for apparently conflicting 'should'-judgments to arise in cases like these, too. But since it's clearest (and least controversial) in the case of beliefs and intentions, I'll continue to focus on them in what follows.

It's worth noting that *actual* failures of "fit" between attitudes are inessential in generating the puzzle. For it turns out that apparently conflicting rational 'should'-judgments will usually arise *whenever* a subject has a "bad" attitude—i.e. one that is unreasonable or unjustified—regardless of whether or not it actually engenders incoherence. And that's because bad attitudes, when combined with one's other attitudes (which may or may not themselves be bad), will invariably have various "downstream" effects, at least some of which will also be bad. In particular, whenever one has a bad attitude—call it BAD—there will be other attitudes that one is rationally required (or, more generally, pressured) to adopt or revise in virtue of having BAD, and at least some of those attitudes will also be bad in the sense of not being adequately supported by one's evidence or reasons. Here as elsewhere, errors tend to propagate.

---

[4] It's an open question how different kinds of attitudes are related to each other. It's *prima facie* plausible to take (most of) them to be distinct, but that's compatible with there being intimate constitutive and/or normative relationships between them.

We can appreciate this point by considering Tom's coherent twin, Tim. Like Tom, Tim believes that he is Superman and that Superman can fly, but unlike Tom, Tim hasn't seriously considered the question of whether he himself can fly—he simply hasn't (yet) put two and two together. Suppose, however, that upon considering the question Tim goes ahead and makes the obvious inference, thereby coming to believe that he can fly. In such a case there's a fairly straightforward sense in which Tim *believes as he should*—it's the doxastic response that's "called for" by his standing beliefs—even though, of course, there's another sense in which Tim does *not* believe as he should, since believing that he can fly runs contrary to all his evidence. For another (and slightly more realistic) example, consider Nina, who is convinced the earth is less than 10,000 years old, and has no conflicting attitudes. Though aware that the scientific community says otherwise, she (wrongly and unjustifiably) dismisses them as being part of some elaborate and sinister conspiracy. She then hears reports of a previously unknown dinosaur bone being discovered nearby. How old should she think the dinosaur bone is? Here again I think we are (or at least can be) pulled in different directions. In one sense, it seems clear that Nina should believe that the bone is less than 10,000 years old—after all, the dinosaur bone couldn't be older than the Earth itself —but in another sense it seems equally clear she shouldn't. Instead, she should believe that the bones are much older.

The same sort of cases arise in the practical realm. Suppose Jane has a sister, Joan, who is a fastidious bookkeeper and fiercely competitive. So, naturally enough, upon learning of Jane's intention to count the blades of grass in the garden, Joan decides to do the same—and to finish the count before Jane does, keeping careful track of the blades she's counted along the way. In such a case, although Joan has adopted the means necessary to achieve her end we're still faced with the question of whether she *should* take the means—is that the rational thing to do? And here again I think we're of two minds. There's a

sense in which she should—how else is she going to finish before Jane?—and another sense in which she shouldn't—her time is clearly better spent doing something else.

Actual incoherence on the part of the subject is thus inessential to the puzzle. The lack of coherence in the initial examples involving Tom and Jane, just like their extremity, merely help serve the heuristic purpose of bringing the different intuitions into particularly sharp relief.

## 1.2 A solution: two dimensions of rationality

There are a variety of possible responses to, and diagnoses of, the puzzle just considered. One option, which I favor, is to take appearances at face value: there appear to be two distinct dimensions of rational evaluation because there *are* two distinct dimensions. In one sense, rationality is a matter of correctly responding to reasons—or so the popular slogan goes.[5] This is the sense in which it is rational for you to believe that the sun will rise tomorrow and 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. (I take talk of evidence to roughly approximate talk of good epistemic reason(s). More on this later.) Likewise, it is rational to exercise regularly as well as look both ways before crossing a busy 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 *reasons rationality*.[6]

---

[5] At least one's epistemic and practical reasons. Perhaps there are other kinds of reasons (e.g. moral) that are less directly relevant to rationality.

[6] I should note that there are several distinct strains in our thought and talk about 'reason(s)' as well, only one of which is of present concern. I'll address some initial complexities later in this chapter, and others in Chapters 6 and 7.

In another sense, however, rationality is a matter of coherence, or having the right structure hold among one's mental states or attitudes, independently of whether those attitudes are reasonable or justified. This is the sense in which it is rational for you 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 is irrational to believe something you think is unsupported by the evidence or do something you believe you shouldn't do. The relevant notion of coherence is thus a broad one, and a 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 any logical inconsistency in contents. Call this *structural rationality*.[7]

Distinguishing between reasons rationality and structural rationality provides a straightforward solution to the puzzle above: the apparently conflicting 'should'-judgments arising in cases involving bad attitudes are not *actually* conflicting. They simply reveal different dimensions of evaluation that we're sensitive to, and in that respect should thus be no more puzzling than those generated by the competing demands of morality and self-interest.[8]

The distinction between reasons rationality and structural rationality is at least latent in the writings of various philosophers, though there's little consensus on how exactly it is to be drawn or how significant it is supposed to be. Some bestow the honorific title of 'rationality' on just structural rationality, opting for another label to denote reasons rationality, while others prefer the reverse. I, on the other hand, think it's best to mark the distinction as one between two dimensions of rationality,

---

[7] Cf. Scanlon (2007), who draws a related, but also importantly different, distinction.

[8] This is not to say that they aren't puzzling. Among other things, difficult questions remain concerning, how (if at all) the different verdicts aggregate, and how, if at all, the difference dimensions relate to each other.

where rationality *simpliciter* is a function of both. On this view, to be fully rational is to have a fully justified, coherent set of attitudes.

However, even if my way of marking the distinction is taken to be mere terminological stipulation, it still has important methodological upshot. For given that our pre-theoretic use of 'rational' and its cognates fails to reliably discriminate between facts about coherence and facts about reasonableness, and given that there's a need to distinguish the two—a religious fundamentalist or conspiracy theorist, for instance, might have a set of beliefs that is quite coherent yet far from justified—it's useful to adopt terminology that demands unqualified judgments concerning "rationality" be disambiguated. Doing so will put us in a better position to accurately handle our (and others') otherwise slippery judgments concerning what the "rational" response is, or would be, in a given situation, as well as what the ingredients are that determine the answer. Ultimately, though, what matters is not the terminology used to demarcate the dimensions, but rather recognition of and sensitivity to the dimensions themselves.

It's worth emphasizing from the outset, however, that talk of reasons is not forced on us. The central contrast I'm concerned with could instead be framed, for instance, as one between coherence and justification, insofar as justification (like reasons) is plausibly thought to come in different flavors—epistemic, practical, moral, etc—and isn't merely a matter of coherence between attitudes.[9] Alternatively, the contrast in the theoretical case (e.g. Tom) could be framed as one between what coherence requires and what one's evidence supports, but because evidence is an inherently epistemic notion we would need some other normative notion—such as reasons or justification—to capture the contrast in the

[9] This is a substantive, but plausible, assumption. "Pure" coherentists about epistemic justification will of course demur. I'll simply be assuming the pure coherentism about justification is false; this is compatible with some form of "hybrid" or "impure" coherentism being true (cf. Berker (forthcoming)).

practical case (e.g. Jane). For the sake of uniformity, then, I'll generally talk in terms of reason(s) and justification, and only occasionally indulge in talk of evidence when discussing theoretical rationality. Choice of normative ideology is of course important—a point this dissertation is intended to illustrate at length—and I'll be subjecting reason(s)-talk to critical scrutiny in Chapters 6 and 7.[10] But I'm hopeful that many of the most important points I'm concerned to make, both substantive and methodological, are translatable into other ideological frameworks (though perhaps not without remainder).

Ultimately, of course, what matters is not terminology *per se*, but instead sensitivity to the different dimensions of evaluation themselves. It's our sensitivity to the latter that should inform our ideology, rather than vice versa. But we have to start somewhere, and given how widespread talk of reason(s) and (to a lesser extent) justification is, and how natural such talk connects up with talk of

---

[10] Here's a quick and dirty sketch of how I think talk of 'reason(s)' (in its normative sense), 'justification', and 'evidence' relate. First: as I'll argue in Chapter 6, reasons (count noun) are ultimately to be understood in terms of reason (mass noun). In slogan form: reasons are "sources" of reason. The mass noun in turn tracks facts about normative support—to say there is (a certain amount of) reason to A is just to say A-ing is normatively supported (to a certain degree)—and so it follows that reasons are sources of normative support. What kind of normative support? It depends. There are lots of different kinds of normative support (or reason), and hence many different "kinds" of reasons—moral, prudential, epistemic, and the so on. (As we'll see, not everything that intuitively counts as a reason is plausibly relevant to rationality—far from it. Which ones do and which ones don't is of course a matter of controversy; I'll address the question, in part, below.) Justification is thus akin to reason (mass noun), since it too denotes a generic kind of normative "stuff"—or, more accurately, a generic kind of normative "force". Why a kind of force, and not just a kind of stuff? Because forces, unlike stuff in general, have both *magnitude* and *direction*, and these are defining features of both reason and justification. That is to say: both reason and justification comes in degrees—you can have more or less reason/justification to do something—and in that sense have magnitude, and both are inherently *directional*—reason/justification is always reason/justification *to* do something, or *not* to do something, and hence always plays either a supporting or opposing role. They're also inherently *relational*—reason/justification is always reason/justification *for* someone to do something, or for them not do to something, and hence always plays either a supporting or opposing role relative to a (possibly implicit) agent. Finally, there are at least two importantly different (though not completely unrelated) strains in our ordinary thought and talk of evidence—in one sense, evidence *of p* is something that is reliable sign or symptom of *p*, while in another evidence *concerning p* is that which justifies, or (more generally) provides support for, believing *p* (cf. Kelly (2014)). In the latter sense, evidence plays the same role as epistemic reasons do—both are functionally defined as sources of epistemic (and hence normative) support.

rationality, it's as good a place to start as any. The central contrast between dimensions of broadly rational evaluation I'm concerned with is one that I think everyone needs to recognize in one way or another. This includes various forms of both "internalism" and "externalism" about reasons, evidence, and/or justification. I myself have broadly internalist sympathies, yet everyone should be able to recognize the difference between a justified belief and a merely internally coherent one, as well as our tendency to use 'rational' to characterize the latter (as well as the former, at least to some extent). Even the arch-externalist Goldman (1986), for instance, is careful to distinguish between epistemic 'justification' and 'rationality', and his reliabilist theory is offered as an account of the former, not the latter. But there's nothing preventing someone like Goldman from recognizing more than one dimension of "rational" evaluation. The same goes for other types of externalists. And there may be some gain—for example, it would allow them to capture the sense in which unwitting victims of mass deception whose attitudes are in internal harmony as well as appropriately responsive to their experience are "doing well" from an epistemic point of view (even if they're not doing *as* well as their undeceived counterparts), whereas those victims whose attitudes are *not* appropriately responsive to their experience are *not* doing well, even if they manage to maintain internal coherence between their attitudes. To admit two "internal" dimensions of broadly rational evaluation is thus perfectly compatible with recognizing other, more "external" forms of evaluation, as well as with prioritizing the latter. This point applies to knowledge-first approaches to epistemic normativity like of Williamson (2000) as much as it does to reliablists like Goldman.[11]

Even though there's room for greater plurality, then, in what follows I'll proceeding as if there are just two dimensions of broadly rational evaluation, with structural rationality solely being a matter of

[11] Indeed, Williamson himself has started moving in this direction—see his (forthcoming).

how things stand mentally—and in that sense internally—with the subject and reasons rationality being at least somewhat constrained, if not entirely determined, by the subject's mental state (or "perspective" more generally). Whether we should be externalists of various kinds about evidence, reasons, justification, and the like, is thus a further question that I'll remain for the most part neutral on.

It's worth emphasizing upfront that I'm *not* assuming that facts about coherence are transparent or "luminous" to one, even upon reflection. Our introspective judgments are highly fallible, and we can be wrong or misled about our attitudes just like we can be wrong or misled about factual matters in general, including facts about our reasons and what they support. We may generally have *better* (and more "immediate") access to facts about our attitudes than we do about facts about the external world, but it doesn't follow that we have *perfect* or even *highly reliable* access. So although facts about coherence are internal in the sense of having to do with (relations between) one's mental states or attitudes, they needn't be internal in the sense of having immediate introspective access to them.

Although the form of pluralism I'm adopting is relatively modest and intended to be as ecumenical as possible, it remains controversial. Indeed, many are simply insensitive to the apparent distinction between reasons-responsiveness and coherence, and proceed on the default assumption that our use of 'rational' and its cognates is univocal. But even among those who *are* sensitive to it, not everyone takes it to be theoretically significant. Some—call them *deniers*—argue that, contrary to appearances, only one dimension of rational evaluation is genuine, or genuinely significant. Whereas some deny the rational significance of coherence as such, and hence deny structural rationality, others deny the rational significance of reasons as such (though not their significance *simpliciter*), and hence deny reasons rationality. Still others—call them *unifiers*—offer theories that can be seen as attempting to provide a single, unified account of our judgments of irrationality.

## 1.3 A sociological observation

It's an interesting—and to my mind quite striking—fact that although the distinction between reasons rationality and structural rationality cross-cuts the practical/theoretical divide, and although practitioners on both sides of the divide have (at least to some extent) recognized the distinction, the dominant *reaction* to the distinction by those on the practical side has been very different from those on the theoretical side. For example, the dominant tendency among epistemologists has been to focus on, and emphasize the importance of, responding correctly to one's *evidence* (which I take to be roughy equivalent one's epistemic reasons, since both evidence and epistemic reasons play the same role of providing epistemic support), where this involves not just believing what the evidence on balance supports, but also doing so "in the right way" or on an appropriate basis. The default assumption oftentimes seem to be that doing so is both necessary *and sufficient* for being epistemically rational.[12] More generally, the epistemological literature has focused on the nature, variety, and sources of epistemic *justification*, and it is almost universally assumed that justification is *not* a matter of coherence.[13] Coherence as such is then typically ignored or dismissed as unimportant. And this is a shame—structural rationality (or whatever we call it) is philosophically and explanatorily important, distinctive, and its verdict can conflict with that of reasons rationality.[14]

---

[12] [Citations: Cohen, Kelly, Feldman and Conee, Hieronymi, etc etc.]

[13] Or at least not *merely* a matter of coherence. So-called "impure" or "non-doxastic" forms of coherentism allow non-doxastic states such as experiences to also play a role. See Kvanvig and Riggs (1992), Pryor (2005), and Berker (forthcoming) for more on this hybrid view.

[14] One way structural rationality is obviously important is in the prediction and explanation of human action, and so is relevant not just within philosophy but also in the various social sciences, such as psychology and economics. Indeed, it seems to be a precondition of being a rational, intentional agent in the first place—of being a believer and an intender—that one is at least *generally* disposed to be structurally rational. This is a central theme of Davidson's work on rationality (collected in his (2004)), among others.

Whether or not it's a truism, it's hard to deny that slogans like "rationality requires one to respond correctly to one's evidence" are latching onto *something* important. Maybe rationality requires more than that (I think it does); maybe there are other things of epistemic interest than rationality (I think there are). What's striking is that whereas there seems to be little doubt that evidence, or epistemic reasons more generally, are of rational significance, the *practical* version of the reasons-responsiveness view is very much controversial. Indeed, if anything the dominant assumption seems to be that rationality is a matter of coherence, or internal harmony—that it is first and foremost (and perhaps exclusively) a matter of one's attitudes relating to each other in the right sort of way, rather than to one's experiences or the world itself.[15] So there's much more of a debate in the literature on practical rationality between the relative merits of the reasons-responsiveness and coherentist views.[16] This trend has been in place for at least the past two decades, and yet has only recently been remarked upon.[17]

Again, there are a variety of different responses to the apparent duality in our ordinary thought and talk about rationality. Linguistically, I favor a rather deflationary response. Philosophically, I favor a form of pluralism. On the linguistic side, I think it's rather clear that 'rationality' and its cognates can be used refer to reasons-responsiveness as well to coherence, and so are at least polysemous.[18] What's more, I doubt that one use is substantially more significant or "central" than the other. Both parties to

---

[15] [Citations: Broome, Darwall, Davidson, Kolodny, Setiya, Scanlon, Wallace, Way, Ridge, Smith, etc.] Broome himself focuses primarily on practical cases in arguing against the reasons-responsiveness view, and even admits (though does not play up the fact that) there are evidential constraints on belief—that theoretical rationality is at least partly a matter of responding correctly to one's evidence.

[16] The difference between the two dimensions of rationality is obviously important when it comes to "rational" decision theory, and what the relevant "value" and "probability" functions are supposed to represent (actual or reasonable preferences? actual or reasonable credences?), though they are oftentimes conflated or mixed up. [Citations? Quotes?] These differences matter.

[17] See, for example, Lord (2013a, 2014a), Sylvan (2014, ms), and Worsnip (forthcoming).

[18] Cf. Buchak (2014), Kolodny (2005).

the debate over the nature of rationality are latching onto something important. Each is just latching onto something different than the other is. It's a mistake to think that there's this one thing—rationality—that is *either* solely a matter of reasons-responsiveness *or* a matter of internal coherence. If anything, rationality *simpliciter* is a matter of both. There's (what I'm calling) reasons rationality and there's structural rationality. I think they both deserve the name 'rationality' as much as the other, and in Chapter 5 I'll suggest there's an important way in which they're related that justifies thinking of them as two dimensions of rationality rather than two completely unrelated dimensions of evaluation, as (say) morality and athletic prowess are.[19] My main reason for being a pluralist is that reasons rationality and structural rationality come apart in both directions, and so neither guarantees the other. (I'll argue for this in 1.5.) Focusing on one to the exclusion of the other will be ignoring something important.

I realize there are some potential drawbacks to being a deflationary pluralist. For example, it turns out that many (though certainly not all) of the debates between proponents of reasons-responsiveness and coherentist views of rationality are merely verbal. This sort of dialectical situation is not uncommon within philosophy, and its recognition is a form of progress, even though it may not be the kind of progress we value most.[20] Yet even if some of the debates between them are merely verbal, much of substance remains, including the articulation and defense of the respective (non-competing) views.

---

[19] In this way, although I'm an advocate of what Worsnip (xx) calls the "disambiguating response", I *don't* think it's just a linguistic or "sociological quirk" that 'rational(ity)' has been used both ways.

[20] This situation arguably arises at least to *some* extent in most major debates. See, for example, Fischer and Tognazzini (2011) on the need to distinguish a variety of different "kinds" of moral responsibility and the relevance of such distinctions to various related debates, as well as Zimmerman (2009), who suggests his seemingly substantive disagreement with Sher (2008) over the epistemic condition(s) of morality responsibility is at least partly verbal.

So although I don't think there's an interesting debate *between* the two views, a lot of interesting and difficult issues remain. One of the most pressing issues, for example, concerns the nature of each dimension, and how they're related. These questions will be the main focus of this dissertation. For it has been standardly assumed that structural rationality is fundamentally requirements-based—that there are a distinctive set of requirements such that to be coherent is to satisfy (or at least not violate) those requirements. I think that's a mistake. The nature of structural rationality is much more akin to the orthodox "pressure-based" conception of the nature of reasons rationality. (I'll explain what this amounts to in due course.) And it has also commonly been assumed that our use of 'reason(s)'—in its ordinary normative sense—is a reliable guide to the things that ultimately matter when it comes to reasons rationality. This too is a mistake. Our use of 'reason(s)'—like our use of 'cause(s)' and 'explanation(s)'—exhibits a kind of disciplined promiscuity, and as a result shouldn't be relied on when engaging in substantive normative inquiry. Or so I'll argue.

Since clear-eyed denial and unification represent important alternatives to the rationality-pluralism I favor, however, the rest of this chapter will be concerned to argue against them. Doing so will set the stage for the main task of the dissertation.

## 1.4 Structural deniers

I'll begin by considering those who deny the rational significance of coherence as such, and hence of structural rationality. Call them *structural deniers*. Structural deniers deny that there's anything *distinctively wrong* with having incoherent attitudes. (Reminder: here and throughout I'm using 'attitude' in an artificially broad sense to include *absences* of attitudes.) However, structural deniers needn't, and typically don't, deny that there's *something* wrong with those who have incoherent attitudes—they just deny

there's something wrong with incoherent agents *because* they're incoherent. They thus recognize the need to explain (away) intuitions that something has gone wrong in cases of incoherence, and face the explanatory burden to accounting for intuitions of wrongness in terms of the violation of other (genuine) norms, and in particular those involving reasons.

Kolodny (2007b, 2008a, 2008b), for instance, is concerned to argue that although incoherence isn't itself problematic, it's a reliable indicator of something that *is* problematic.[21] For, as he argues, there are at least many cases in which having an incoherent combination of attitudes guarantees that one (or more) of those attitudes is unjustified or unreasonable, and hence an attitude that one shouldn't have—independently of the attitudinal conflict. This is easiest to see in cases of involving inconsistency between beliefs. Kolodny agrees, for example, that it's irrational to believe *p* and to believe not-*p*. But he denies that that's because believing *p* and believing not-*p* violates the demands of structural rationality. Instead, it's because if you believe both *p* and not-*p*, you are guaranteed to believe something you lack adequate reason to believe. And that's because it's impossible to have adequate reason to believe both *p* and *not-p*. So at least some forms of incoherence seem to guarantee unreasonableness.

According to (what I'll call) the *wholesale* structural denier, this pattern generalizes—having an incoherent combination of attitudes guarantees that you either have an attitude that, given your reasons,

---

[21] Kolodny sees himself as defending and generalizing a view first suggested by Raz (2005). It's worth noting up front that Kolodny is something of a moving target. In his (2005), for instance, Kolodny indicates an openness to accepting coherence requirements as merely "evaluative" requirements—necessary conditions for qualifying for a certain kind of appraisal—and the success of his "Transparency Account" depends on the truth of a small handful of "higher order" coherence requirements governing one's attitudes and one's beliefs about their normative status. Broome would disagree with the "mere necessary" condition characterization of coherence requirements in general, but otherwise there might not be much difference. In his subsequent work, however, Kolodny is skeptical of even that weaker claim. Nonetheless, he doesn't really *argue* against it, and relies on overly simplistic divide between "normative/deontic" standards like morality and "merely evaluative" standards involved in (e.g.) chess and beauty. I think structural rationality corresponds to a distinctive type of normative criticism/assessment that is more robust than that associated with chess or beauty, and yet not on a par with (say) morality. I'll be glossing over these issues.

you shouldn't have or else lack an attitude that you should have.[22,23] Instead of positing two genuine, and genuinely distinct, dimensions of normative evaluation in an effort to explain our intuitive judgments about (ir)rationality, we would only need one—reasons rationality. That is to say, if facts about reasons—and, in particular, facts about which possible combinations of attitudes can be adequately supported by reasons—were to guarantee that whenever an agent has an incoherent combination of attitudes she also has an attitude that is unjustified, then we would seem to be in the enviable position of being able to accommodate the intuitive *verdicts* of structural rationality without accepting the *reality* of structural rationality itself (at least insofar as it is supposed to diverge from reasons-responsiveness).[24] Our sense that something is wrong with agents who are incoherent would be traceable to the fact that they fail to respond correctly to their reasons. Reasons rationality might thus seem to have all the resources needed to satisfactorily explain the intuitive verdicts about particular cases.

## 1.5 Against structural deniers: reasons-responsiveness without coherence

---

[22] One of course might be a *selective* structural denier rather than a *wholesale* one—that is, one might deny that some kinds of incoherence are (themselves) rationally problematic without denying that *all* kinds of incoherence are. I'll be concerned with wholesale structural deniers.

[23] Or, more modestly, what matters for the structural denier is that *intuitively problematic* incoherence entails unreasonableness. Perhaps there are cases in which apparent cases of incoherence are not intuitively problematic, or irrational in any sense (cf. Kolodny (2007, 2008)). The structural denier is concerned to deny that there's anything *distinctively wrong* with being incoherent; if there are cases in which there's *nothing* wrong with being incoherent, so much the better (for them). However, I'm working with a notion of incoherence according to which it is (by definition) intuitively problematic—to be incoherent (in the sense I'm interested in) is to be negatively appraisable in a certain distinctive way. (Contrast Foley's (19993) use of 'incoherent' to mean 'probabilistic incoherence', which he doesn't find inherently problematic—and neither do I.)

[24] Compare attempts to explain the plausibility or acceptability of various moral claims while denying the reality of moral facts. [Citations?]

There are at least two major problems facing wholesale structural deniers.[25] The first problem is that the proposed explanation of what's wrong with having inconsistent beliefs does not, in fact, generalize. This isn't terribly surprising given how high the bar is for success: the wholesale structural denier has to make it plausible that for *every* incoherent set of attitudes *S* (where this may include the *lack* of an attitude), there are facts about reasons and how they combine that entail there is something wrong with *S*. The problem is that there seem to be a variety of cases in which it is possible to have individually justified but jointly incoherent attitudes. These will then be cases in which there's something intuitively wrong with a certain combination of attitudes, and yet facts about reasons are silent as to why. Even though there's controversy concerning each of the cases I'll consider, taken together they pose a serious challenge for the (wholesale) structural denier. They also help make plausible the main contention of this chapter—namely, that there are two genuinely distinct dimensions of rational evaluation, and hence two importantly different normative phenomena that need to be kept distinct in our theorizing, with neither being reducible to the other.

The possibility of a set of attitudes being reasons rational (reasonable) while failing to be structurally rational (coherent) is most obvious in the practical realm. To use a stock example: suppose that Buridan's ass is stuck between two equally attractive bales of hale and intends to eat each, despite knowing he can only eat one. Here the problem doesn't seem to be with any of the attitudes on their own, each of which is justified. Instead, the problem seems to be with the incoherent combination of attitudes itself. Or suppose Buridan's ass intends to eat the bale of hay on the left but is too lazy to do what he knows is necessary for achieving his end—namely, to start moving to the left. Here again the

[25] Both problems have been noted before and discussed fairly extensively. See, in particular, Way (2014, forthcoming).

individual attitudes seem unproblematic on their own, given that each is adequately supported by his reasons.[26] But their combination results in a paradigmatic case of means-end incoherence.

Any case in which there are multiple incompatible options or attitudes, each of which is equally well-supported by one's reasons, can be used to make the same point.[27] These sorts of cases, as well as the the general problems they pose, have received a lot of attention.[28] But as Way (2014, forthcoming) notes, these aren't the only sorts of cases from the practical realm that pose a challenge for the wholesale structural denier. Cases involving incommensurability and supererogation do as well (assuming there are such). These, too, are cases in which there is more than one alternative or attitude that is adequately supported by one's reasons, and hence permitted by reasons rationality, yet actually adopting more than one of the relevant attitudes would result in incoherence. For our purposes we can think of cases of incommensurability as involve two or more alternatives, neither of which is better supported by the totality of one's reasons and yet which aren't equally well-supported either.[29] This is evidenced by the fact that a small increase in the amount of support enjoyed by one of the alternatives would not "break the tie" and make it better supported than the other(s)—the alternatives would remain "on a par", despite the increase in support enjoyed by one of them. As Way points out, for many people there may be incommensurable reasons in this sense "to go to law school or graduate school in philosophy, to visit Salisbury Cathedral or Stonehenge, to listen to the Beach Boys or the Beatles, or for Sartre's famous

---

[26] This includes his failure to intend to move left, since he has sufficient reason to move right, and hence not move left. At least this is true given a suitable "transmission" principle.

[27] And, as Bratman (1987) emphasizes, such cases are legion in ordinary day-to-day life. [Page/quote?]

[28] For this point or the more general problem described below, see, e.g. Batman (1987), Kolodny (2007), (2008), Ross (2012); Schroeder (2009), Wedgwood (2011), Way (2012).

[29] Cf. Raz (1986) and Broome (2000), among many others. As Chang (1997) emphasizes, there's actually a cluster of related but distinct phenomena that often get conflated. Like Way, I don't intend to be taking a stand on whether the cases I'm interested in best thought of as cases of incommensurability as opposed to (e.g.) incomparability, parity, or rough equality.

student to stay home with his mother or to fight for the resistance" (page citation). In all these cases, it seems possible to have inconsistent intentions or be means-end incoherent without having an attitude that, as far as reasons rationality is concerned, you shouldn't have, or failing to have some attitude that you should.

The same goes for cases of supererogation—i.e. cases where there are multiple incompatible options, each of which is *adequately* supported but (at least) one of which is *better* supported than the rest—these are more controversial.[30] The most compelling cases of supererogation are moral ones—e.g. sacrificing your life in order to save the lives of others may be morally admirable, yet arguably isn't required—but it's not implausible to think they arise in the practical domain as well.[31] Perhaps the choice between becoming a truly excellent pianist and becoming a skilled (but not excellent) guitarist, where you are confident you can become one of the two but not both, is an example.

Whether analogous cases involving "mere permissibility" arise in the theoretical domain is more contentious. One kind of case hinges on whether there is ever more than one doxastic attitude (partial or full) it would be reasonable—and hence permissible—to adopt, given a particular body of evidence.[32] Assuming, as seems plausible, there are at least *some* cases in which reasons rationality gives you at least *some* leeway concerning which doxastic attitude to adopt, then we can construct cases in which an agent has a set of individually reasonable but jointly incoherent doxastic attitudes. Kelly (2014) provides a plausible candidate for a permissive case:

> Suppose that six months before the U.S. presidential election, it's quite unclear whether the Democratic or the Republican nominee will win [and there is no third candidate]…I possess a

---

[30] Again, I'm glossing over a range of issues. See [Citations].

[31] The practical/moral divide is of course somewhat of a murky one.

[32] See Kelly (2014) and White (2014) for a recent exchange on this topic. Like Kelly, I suspect that there relatively few "permissive" cases, and that most of them aren't all that permissive.

> large body of information [bearing] on this question [and it is on balance] somewhat more likely that the Democrat will win than not…Suppose [that] you and I agree on the basis of our common evidence that the Democrat is more likely than not to be elected [but that] it's far from a sure thing. The only difference between us is this: you're a bit more cautious about the Democrat's prospects, and so give a bit less credence to the proposition that the Democrat will win than I do. (299-300)

As Kelly notes, "the natural verdict about the case is that it's consistent with everything that's been stipulated so far that you and I might both be fully reasonable in our opinions about the election, despite the fact that those opinions are not identical" (300). Notice, though, that this opens up the possibility of an agent having a set of attitudes that is justified but not coherent. For suppose that even though you give a bit less credence to the proposition that the Democrat will win than I do, you give virtually the same amount of credence to the proposition that the *Republican* will win as I do—and this despite you (and I) knowing that one and only one of the two candidates will win. Indulging in idealization, we may suppose your credence in the former is (roughly) .53 whereas your credence in the latter is .44, despite your credence in their disjunction being extremely close to 1.[33] We may suppose this subtle incoherence in your credal state would manifest itself in your betting behavior, though it in fact eludes you since you don't seriously consider placing any bets.[34] The important point is just that since the credence I give to the proposition that the Republican will win is reasonable (or so we're assuming), and since we share the same evidence concerning the outcome of the election, the credence you give to the proposition is reasonable as well. Here as elsewhere a lot depends on the details, but if a case like this is possible then we should be open to the possibility of there being justified (though unrecognized) incoherence among

[33] I say "extremely close" since you may allow for the possibility of some cataclysmic event preventing the election from taking place, though you consider it exceedingly unlikely.

[34] We may also suppose the incoherence isn't obvious upon introspection—our judgments concerning our own dispositional mental states are highly fallible, after all. But just because the incoherence eludes you doesn't mean the incoherence isn't there.

our doxastic attitudes, at least with respect to graded attitudes like credences (as opposed to outright attitudes like beliefs).

Another kind of case that illustrates the possibility of conflict between reasons rationality and structural rationality—and one that I find particularly interesting—involves the potential impact of higher-order evidence, which may or may not be misleading. These sorts of cases arise in both the practical and theoretical domain. Suppose, for example, that I have excellent, perhaps even decisive, reason to (intend to) do something—say, finish my dissertation. And yet suppose a trusted advisor with an otherwise stellar track record of good advice tells me otherwise—she insists that it would be best for me to do something else, such as to go on vacation. We may suppose she's perfectly sincere, despite being mistaken. In such a case, what should I do—what would be most reasonable? Given how I've set things up, the answer seems clear: finish my dissertation. Although it's plausible the advisor's testimony has at least *some* effect on the (first-order, practical) question of what I should do, making it less reasonable to finish my dissertation than it would have been otherwise, I see no reason to suppose that it has a *swamping* effect, outweighing all the first-order considerations that so strongly favor finishing my dissertation, or even that it has a *neutralizing* effect, balancing the scales so as to make it equally reasonable to (say) go on vacation. At the very least I see no reason to suppose that it is *guaranteed* to have such effects, no matter how we spell out the details of the case and no matter how strong we stipulate the first-order reasons to be. But this doesn't settle the question of what the *epistemic* effect the advisor's testimony is, or should be—that is, its effect on the question of what it would be most reasonable for me to *believe* I should to do. Since she's a trusted and (up to this point) highly reliable advisor, her testimony clearly carries significant evidential weight—and enough, we may suppose, to make it reasonable for me to believe I should go on vacation.

The first-order, practical question of what I should do—in the sense of what I have most reason to do, or what it would be most reasonable to do—is thus a separate question from the higher-order, theoretical question of what I have most reason to *believe* I should do, and the answers may very well fail to "line up". So although I realize a lot depends on the details, I nonetheless think it's plausible that there can be cases in which the most reasonable thing for me to do, all things considered, is one thing (e.g. finish my dissertation) and yet the most reasonable thing for me to *believe* I should do is another (e.g. go on vacation). If I were to then adopt the attitudes it would be most reasonable for me to adopt, I would end up intending to do something I believed I shouldn't do, and hence be akratic—a paradigm case of incoherence.

Similar cases arguably arise in the theoretical realm as well. Just as one might have misleading higher-order evidence concerning what one's first-order reasons (on balance) support *doing*, so one might have misleading higher-order evidence concerning what one's first-order reasons, or evidence, support *believing*. Perhaps the higher-order evidence comes from (at least apparent) expert testimony, or perhaps it's furnished by past experience. There are a variety of such sources. And so just as it seems possible to have a justified but false belief concerning the first-order question of what one should do, so it seems possible (in principle) to have a justified but false belief concerning the first-order question of what one should believe. Indeed, this possibility has been explored and defended) by a number of recent authors, each of which have offered examples that are good candidates for being cases in which a subject is justified in believing that they're not justified in believing *p* even though they *are* justified in believing *p,* or in which they're justified in believing that they're justified in believing *p* even though they are *not*

justified in believing $p$.[35] To "follow the evidence" and adopt the beliefs they're justified in having would thus result in believing something they believe they (epistemically) shouldn't believe, or in failing to believe something they believe they should believe, and thus be epistemically akratic—another important form of "inter-level" incoherence.

One interesting feature of cases involving misleading higher-order evidence is that it's (at least usually) impossible for one to know, or stably believe, that one is in such a case. [FILL IN] There are other (though, again, controversial) examples that suggest something even stronger—namely, that it's possible to have a set of justified yet *recognizably* incoherent set of attitudes.

Warren Quinn's (1990) "puzzle of the self-torturer", for instance, provides a *prima facie* plausible example of a subject who is fully justified in having a set of recognizably intransitive preferences. He considers a person with a special electric device attached to him. It has 1,001 settings, ranging from 0 (off) to 1,000 (max). Each increase in setting leads to a tiny increase in amount of electrical current applied to his body, and hence a negligible (and phenomenologically indistinguishable) increase in physical pain. Each week the "self-torturer" is given a choice to either stay put or else advance the dial one setting, but "he may advance only one step each week, and he may never retreat" (79). For each advance he gets a large sum of money (say, $100,000). Although the self-torturer cares about his overall level of physical comfort he also cares a great deal about a variety of things that large sums of money can buy. And since he doesn't (indeed, can't) notice any difference in comfort between adjacent settings, for any two successive settings $n$ and $n+1$, he prefers stopping at $n+1$ to stopping at $n$. However, he realizes that were he to continue to advance each week he would "eventually reach settings that will be so

[35] See, for example, Pryor (2004, ms), Greco (2014), Horowitz (2014), Lasonen-Aarnio (ms), Littlejohn (forthcoming), Williamson (2011), and Worsnip (forthcoming). The terminology and details vary, as does the purported upshot, but they all present examples than can be (re-)interpreted along such lines.

painful that he would then gladly relinquish his fortune and return to 0" (79). And so he prefers stopping at a low setting (e.g. 0) to stopping at a high setting (e.g. 1,000). As a result, his preferences are intransitive. As Quinn notes, although in such a case it would certainly be irrational for the agent to (intend to) *act* on each preference, it doesn't follow that it would be irrational for him to *have* all of them.[36] On the contrary, Quinn thinks they are "considered and well-informed" and "seem perfectly natural and appropriate given his circumstances", even though "many theorists would condemn his intransitive preferences as irrational" (80).[37]

There's a lot of controversy surrounding what the proper diagnosis of the puzzle, and there's no need for me to enter the fray take a firm stand. Nonetheless, it's worth noting that the debate has largely proceeded on the assumption that "rationality" is a unified phenomenon, and if we give up that assumption then there's an alternative diagnosis available. On this alternative, although it's true that the self-torturer's intransitive preferences are *reasonable* or *justified*, and hence reasons rational (just as Quinn's description of them being "considered and well-informed" suggests), they are nonetheless incoherent considered as a group, and hence are not fully *structurally* rational. If this were the right

---

[36] I'm taking it for granted that intransitivity in one's preferences is a type of incoherence for dialectical purposes, though this assumption is debatable. I myself am not unsympathetic with the thought that self-torturer-style cases show it's possible to be structurally rational or coherent (in the relevant normative sense) while having intransitive preferences—though it isn't structurally rational to intend to *act* on them. Nonetheless, many philosophers seem to think there's something obviously wrong with having intransitive preferences, and this sense of wrongness might be chalked up to the fact that intransitive preferences—at least if recognized—is incompatible with being fully structurally rational, even though it's compatible with being each attitude being fully reasonable.

[37] Although the self-torturer finds himself in a rather far-fetched scenario, Quinn points out that the self-torturer "is not alone in his predicament. Most of us are like him in one way or another. We like to eat but also care about our appearance. Just one more bite will give us pleasure and won't make us look fatter; but very many bites will. And there may be similar connections between puffs of smoke and lung cancer, or between pleasurable moments of idleness and wasted lives" (79).

diagnosis, it would be another way of illustrating how conflicts between reasons rationality and structural rationality might arise.

It's worth noting that although the debate involving the self-torturer has focused almost exclusively on the rational status of such an agent's *preferences,* there's a related (but distinct) question concerning the rational status of such an agent's *intentions*. For although it would be *both* unreasonable *and* incoherent for the agent to intend to advance every week, when we consider each week on its own (rather than as part of a group) it looks like the agent would be justified in intending—and not just *preferring*—to accept the large sum of money that week in exchange for a negligible increase in pain. After all, almost everyone agrees he would be justified in accepting the money at the outset, and the case is set up in such a way that the choice situation each week is (nearly) qualitatively identical to the rest. So if an intention to accept the money is justified the first week, it looks like it would be justified the second week, the third week, the fourth week, and so on. What's more, since intentions can be (and oftentimes are) formed in advance of action, for any given week *w* it looks like the agent he would be justified *at the outset* in forming the intention to accept the money in *w,* at least conditional on being given the choice (since he realizes he might, and eventually will, decline at some point). It's only when such intentions (conditional or otherwise) are considered as a sufficiently large group—i.e. taken collectively rather than individually—that they cease to be justified, and start looking incoherent, given the agent's preference and (presumably) intention to not suffer agonizing pain, which he realizes a sufficiently large group of intentions would guarantee.

The possibility of having a justified but recognizably incoherent set of attitudes arises in the theoretical realm as well. In particular, the large literatures surrounding both the so-called "preface paradox" and the "lottery paradox" contain a host of examples in which it seems reasonable for a subject

to adopt a set of beliefs that is not only inconsistent, but also recognized (or at least recognizable) as such, and so at least *prima facie* incoherent. This claim about incoherence is controversial if we adopt a view of "full" belief according to which it is just a matter of having a (stable) degree of confidence above a certain (perhaps contextually variable) threshold.[38] It is an important and un-argued for presupposition of this dissertation that "partial"/"graded" and "full"/"outright" attitudes are distinct—neither can be reduced to the other—and that both are psychologically and rationally important. This is true not just of belief and its irreducibility to degrees of confidence (perhaps together with some other stuff) but also of intention and its irreducibility to belief-desire complexes (perhaps together with some other stuff). Ultimately this is an empirical claim concerning the states and process belonging to the cognitive architecture of human beings, and there are a range of considerations—philosophical, phenomenological, and empirical—that provide support for it. There are of course other considerations that tell against, but I won't try to adjudicate the issue here.[39] It's worth emphasizing, then, that the assumption of the reality of full attitudes and their irreducibility to partial ones is inessential to the main goals and claims of this dissertation; they could be re-cast exclusively in terms of partial attitudes. Nonetheless, I'll continue to speak of both kinds of attitudes, often with an emphasis on the full attitudes, not just because it helps streamline the discussion but also because I think both are psychologically and philosophically important.

To illustrate the lottery paradox, consider the following example from Schechter (2013):

---

[38] This is the so-called "Lockean" view of belief. It's a form of *reductionism* about full belief—unlike *eliminativism* about full belief, Lockeans accepts the *reality* of full belief, though they think it reduces to, or at least supervenes on, partial belief.

[39] The literature is vast. For some of recent arguments in favor the reality and irreducibility of full doxastic attitudes such as belief and suspension of judgment, see Friedman (2013a, 2013b, forthcoming), Ross and Schroeder (2014), and especially Weisgberg (ms). For some in favor the reality and irreducibility of full practical attitudes such as intention, see Bratman (1987) and Holton (2009).

> Suppose that there is a raffle with 1,000 tickets. One ticket, chosen at random, will win. Suppose that I am aware of this. Suppose that I have a lot of time on my hands, and to fill my time I form beliefs about each of the tickets. In particular, for each ticket, I form the belief that it will lose on the grounds that it has a 999/1000 chance of losing. Presumably I am justified in having each of these beliefs. But if I were to infer from them that all of the tickets will lose, I would not be justified in this new belief. That is because I am aware that some ticket will win. (434)

"Lottery-style" cases like this raise a number of important issues in epistemology[40], but what's most relevant for our purposes is that it illustrates the possibility of having a set of justified, but recognizably inconsistent beliefs—I believe of each ticket that it will lose, and also that at least one of the tickets will win, and the inconsistency between these beliefs is obvious to me. There thus appears to be a residual unhappiness, or structural deficiency, in my total belief state, even though each of my individual beliefs is justified.

As with most of the examples considered so far, there's controversy over the details of the case. Some philosophers, for example, will reject the assumption that I'm justified in believing of each ticket that it will lose. However, doing so would seem to require adopting a kind of general skepticism about merely statistical grounds for belief—one that I (and others) find hard to maintain, especially when it comes to more mundane cases involving chance.[41]

Some of these worries, however, can be sidestepped by focusing on the equally familiar "preface paradox". Abstracting away from the irrelevancies introduced by examples involving authors and their books, Fitelson and Easwaren (2015) offer the following schema for constructing preface-style examples:

---

[40] Schechter, for instance, appeals to it while against the following "closure" principle concerning doxastic justification:

> (**Closure**) Necessarily, if S has justified beliefs in some propositions and comes to believe that q solely on the basis of competently deducing it from those propositions, while retaining justified beliefs in the propositions throughout the deduction, then S has a justified belief that q.

[41] Cf. Schechter (2013). See Hawthorne (2004) for a discussion of some additional complexities.

> Let **B** be the set containing all of *S*'s justified first-order beliefs. Assuming *S* is a suitably interesting inquirer, this set **B** will be a very rich and complex set of judgments. And, because *S* is fallible, it is reasonable to believe that some of *S*'s first-order evidence will (inevitably) be misleading. As a result, it seems reasonable to believe that some beliefs in **B** are false. Indeed, we think *S* herself could be justified in believing this very second-order claim. But, of course, adding this second-order belief to **B** renders *S*'s overall doxastic (full belief) state deductively inconsistent.

Assuming the agent is capable of realizing her full set of beliefs is inconsistent—she needn't to attend to them individually to do so, after all—then case also illustrates the possibility of having a set of justified beliefs that one recognizes to be inconsistent.

The emphasis on the recognizability of inconsistency in such cases is important because it's what helps give them such a strong whiff of "paradox" (or so I think). And that's because *recognized* inconsistency seems to introduce a kind of internal tension or incoherence within one's total belief state that *mere* (i.e. unrecognized) inconsistency does not, and so motivates the thought that these are cases in which reasons rationality and structural rationality come apart.[42] Indeed, it's not obvious that *mere* inconsistency itself counts against the rationality of a belief state at all. As Field (2009) notes,

> [T]he idea that it is always irrational to be inconsistent seems absurd. Indeed, it is natural to suppose that *any* rational person would have believed it impossible to construct a continuous function mapping the unit interval onto the unit square, until Peano came up with a remarkable

[42] I'm open to the idea, however, that preface-style cases show that it's possible to be structurally rational or coherent (in the relevant normative sense) while being recognizably inconsistent. Nonetheless, there does seem to be some residual unhappiness in such cases, and this unhappiness might plausibly be chalked up to the fact that recognized inconsistency is incompatible with being fully structurally rational, even though it's compatible with being fully reasonable. Either way, though, there's another interesting upshot of lottery-style and preface-stylce cases for structural rationality, given plausible normative constraints governing the relationship between full belief and credence—they suggest not only that it's *not irrational,* structurally speaking, to have a deductively closed set of beliefs (a set of beliefs closed under logical consequence), but also that it may be positively *irrational* to do so. See Worsnip (forthcoming-b).

> demonstration of how to do it. The belief that no such function could exist (in the context of certain set-theoretic background beliefs) was eminently rational, but inconsistent. (254)

Although Field himself doesn't distinguish between reasons rationality and structural rationality, I think his example has force interpreted either way. It's clearest in the case of reasons rationality, and this just reinforces the point often made using lottery-style and preface-style cases that one can be justified in having an inconsistent set of beliefs. But it's also *prima facie* plausible in the case of structural rationality as well. To say that many mathematicians, prior to Peano's demonstration, had inconsistent beliefs is one thing; to say that they were structurally irrational, in the (broadly normative) sense of being incoherent, seems to be another.

A similar issue concerning the relationship between logic and structural rationality will arise again in Chapter 2. For now I just want to register the conviction that the *recognition* of inconsistency introduces a certain amount of internal tension or incoherence in one's total belief states in a way that mere inconsistency does not. And this apparent tension, or incoherence, is one that (in cases like those above) facts about reasons or justification are powerless to explain.

I realize some may not share the sense that it is always structurally irrational (to at least some degree) to have a set of beliefs that one recognizes to be inconsistent. They may agree that it's incoherent to believe $p$, believe that $p$ entails $q$, and yet believe that not-$q$, or to believe $p$ and believe $q$ and also believe that either $p$ or $q$ is false. It's at least a little puzzling, however, why when it comes to sets of (full) beliefs, *mere size* makes a difference—that it's structurally irrational to have a smallish set of beliefs one recognizes to be inconsistent but *not* structurally irrational to have a largish set of such

beliefs.[43] It's relatively clear why size matters when it comes to justification. For one can be justified in believing *p* even though there's some (relatively small) chance that *p* is false.[44] And since epistemic risk aggregates, even if one is justified in believe each of a long series of claims (as in the lottery and preface cases), one might not be justified in believing the higher-order claim that all of the first-order claims are true. The likelihood, or "risk", of at least one of them being false is too high.

If we think of full belief as a kind of (doxastic) commitment, however, it doesn't look like an analogous story can be offered in an effort to explain why size of belief set matters to structural rationality, rather than just reasons rationality. Unlike justification, which comes in degrees, commitments are all-or-nothing, and if you're committed to *p* being true and you're committed to *q* being true, it's not clear in what sense you're incurring any additional "risk" relative to your existing doxastic commitments (as opposed to your reasons) in being committed to both being true. Clearly you can be *less confident* that both are true than you are in each individual proposition, and in that sense be less committed to them both being true. And this is important. But insofar as we think of belief as a distinctive doxastic attitude—one that isn't reducible to a certain level or amount of confidence—then we should expect it to make it's own distinctive contribution to structural rationality. And if we do, then it's not clear that having larger sets of beliefs that you recognize to be inconsistent is somehow more coherent, or structurally rational, than having smaller sets of beliefs that you recognize to be inconsistent. Both involve commitments that are recognizably incompatible—in both cases you're *committing* yourself to believing something false—which seems to be a paradigmatic form of incoherence.

---

[43] Analogous points can be made using partial beliefs. It's not usually thought that having a *large* set of probabilistically incoherent credences is any better than having a small one, for example.

[44] And this of course holds more generally: having *sufficient* reason to do or believe something is rarely a matter of having *conclusive* reason to do so.

It's worth comparing (full) belief to the attitude of supposition in this respect. For like belief, supposition is all-or-nothing and is plausibly thought of as a kind of (hypothetical) commitment. And it seems clear in the case of supposition that mere size of the set of propositions is irrelevant to whether such a set of propositions can be coherently supposed. It's just as impossible to coherently suppose that $p_1$, $p_2$, $p_3$,... $p_{1,000}$, and that $\sim(p_1 \wedge p_2 \wedge p_3 \wedge \ldots p_{1,000})$ as it is to suppose that $p_1$, $p_2$, and that $\sim(p_1 \wedge p_2)$. Insofar as beliefs are akin to non-hypothetical suppositions, we should expect the same to hold of them as well.

Even though there's controversy concerning each of the cases I've considered so far, and even though I haven't been able to defend any at length, taken together they make plausible the claim that it's possible to have a set of attitudes that is fully justified without being coherent, and hence that it's possible for structural rationality to come apart from reasons rationality. We thus have good grounds to think that the wholesale structural denier's strategy of explaining the sense in which something has "gone wrong" in cases of incoherence solely in terms of facts about reasons (and how they combine to support different combinations of attitudes) is hopeless. And we thereby have good grounds for taking structural rationality to be a distinct dimension of rational evaluation that is worth theorizing about in its own right.

## 1.6 Against structural deniers: the distinctiveness of (in)coherence

The second major problem for the wholesale structural denier is more flat-footed, and arises even if they're right that incoherence guarantees failure of reasons-responsiveness, and hence even if it's impossible to be reasons rational without being structurally rational. And that's because the structural denier simply doesn't do justice to the phenomena—there seems to be something *distinctively wrong* with

an agent who has incoherent attitudes, and something *distinctively right* with one who has coherent attitudes.

Recall Tom above. In setting up the original puzzle I noted that there seemed to be two things wrong with him, both of which are naturally characterized as rational failings: not only did he have a crazy belief, but he also failed to believe the obvious consequences of other things he believed. The structural denier denies that the second failing is, in fact, a failing. The only genuine failing is the first—the failure to respond correctly to his reasons. But that doesn't seem to do justice to the phenomenon. The inadequacy of such a response can also be brought out by comparing incoherent Tom to his *extra* incoherent twin Ted. Like Tom, Ted believes that he (Ted) is Superman and that Superman can fly, but unlike Tom—who isn't sure whether he himself can fly—Ted is firmly convinced he (Ted) cannot fly. Although both Ted and Tim fail to believe in accordance with their evidence—neither actually has any reason to believe they're Superman—there is a clear sense in which Ted is *even more* irrational than Tom. But this difference isn't a difference in whether their attitudes are justified—in fact, judged by that metric alone, Ted is more *rational*, since he is rightly convinced that he can't fly whereas Tom remains uncertain. So the only explanation for why Ted is worse off than Tom, rationally speaking, is that he's more incoherent.

The second point—that there's something distinctively right about having coherent attitudes—can be brought out by returning to Tom's coherent twin, Tim. Like Tom, Tim believes that he is Superman and that Superman can fly, but unlike Tom, Tim believes that he can fly. Although both Tom and Tim fail to believe in accordance with their evidence—neither actually has any reason to believe they're Superman—there is a clear sense in which Tim is more rational than Tom. But this difference

isn't a difference in whether their beliefs are justified—in fact, judged by that metric alone, *Tom* is more rational, since two of his beliefs accord with the evidence while only one of Tim's does.

The same basic point can be made using any of the other forms of incoherence considered so far, including means-end incoherence, intention inconsistency, and the like. The combined upshot is that we need to appeal to *both* dimensions of rationality to do adequate justice to our (firm, stable) intuitive judgments concerning the full range of cases.[45]

## 1.7 Reasons and rationality

I'll now briefly consider those who, like Broome, grant the rational significance of *coherence* as such but deny the rational significance of *reasons* as such (though they needn't deny the significance *simpliciter* of reasons). Doing so will allow us to begin disentangling some of the different strains in our ordinary reasons-talk, and as a result clarify the nature of reasons rationality. I'll be subjecting reasons-talk to greater—and more critical—scrutiny in Chapters 6 and 7. So what follows should be taken as provisional.

In trying to make sense of ordinary reasons-talk, it's useful (and standard) to begin by distinguishing between the reasons *why* something is the case (so-called "explanatory reasons"), the reasons why—or *for which*—someone does something (so-called "motivating reasons"), and the reasons

[45] It's worth emphasizing that Kolodny himself is very much aware of both problems (though perhaps not their full extent). And in his (2007, 2008a, 2008b) he makes an attempt at addressing both (what he calls) the "violation intuition"—the intuition that something is wrong with incoherent sets of attitudes— and the "satisfaction intuition"—the intuition that something is right with coherent sets of attitudes. His attempt at explaining the latter hasn't received nearly as much attention as his attempt at explaining the former. One reason may simply be its implausibility. See Way (2014, forthcoming) for a convincing rebuttal. Lord (2014) also defends "the myth view" but only addresses the violation intuition—and only provides a error theory for a small handful of coherence requirements. He thinks the explanations generalize, but offers no argument. Perhaps most importantly, however, he says nothing to address the satisfaction intuition.

for someone *to* do something (so-called "normative reasons"). It's that latter that are of present concern —the slogan 'rationality requires one to respond correctly to reasons' is to be understood as expressing the claim that rationality requires one to respond correctly to *normative* reasons. Normative reasons are in turn standardly, if unenlighteningly, characterized as considerations that "count in favor of" or against performing various actions or having certain attitudes.[46]

Standard objections to the claim that rationality requires one to respond correctly to normative reasons take the form of counterexample. For example, Broome (2007a: 167) offers the following "quick objection":

> On some occasion, there might be a reason for you to achieve something but, without any irrationality on your part, you might not believe this reason exists. If you do not believe it exists, then you might well not respond correctly to it, and your failure will not imply any failure of your rationality. Therefore, rationality cannot consist in responding correctly to reasons.[47]

To illustrate, he considers the following case:

> Suppose the fish in front of you contains salmonella. This is a reason for you not to eat it. But there may be no obvious evidence that it contains salmonella. So you might not believe it contains salmonella, and you might eat it, and nevertheless you might be rational. So you are rational even though you do not respond correctly to the reason.

---

[46] I use the fudge term 'considerations' intentionally. I don't intend to take a firm stand on the ontology of normative reasons—on what kinds of "things" (facts, states, propositions, whatever) they are—although I'll often talk as if they're all facts. In Chapter 7, however, I'll provide some grounds for thinking we should be pluralists about the ontology of reasons.

[47] Notice that Broome's target is the claim that rationality *consists* in responding correctly to reasons. I'm only concerned with defending the weaker claim that rationality *requires* one to respond correct to reasons (of a certain sort). But Broome's strategy is to argue against the former by arguing against the latter, since the former entails the latter.

As Broome notes, in this example you are (non-culpably) ignorant of the non-normative fact that the fish contains salmonella, which, we may suppose, constitutes a reason for you not to eat it. Elsewhere Broome also considers a variant in which you realize the fish contains salmonella, but are ignorant of the *normative* fact that it constitutes a reason for you not to eat it.[48] (Perhaps you're simply uninformed about most medical matters, or else have been misled by your otherwise trustworthy doctor.) In either sort case, Broome thinks, one needn't be irrational even though one fails to respond correctly to the relevant normative reasons.

Even though I'm not interested in defending the claim that rationality consists *exclusively* in responding correctly to reasons, I am (provisionally) interested in defending the claim that rationality is *partially* a matter of responding correctly to reasons—and, in particular, that doing so is a necessary condition of being rational *simpliciter*, along with being coherent. And Broome's proposed counterexample, if successful, is a counterexample to both claims.

The problem, however, is that Broome's counterexample relies on an insufficiently nuanced conception of reasons and their relevance to rationality. For within the class of normative reasons there's a further distinction to be drawn. One fairly intuitive—and popular—way of drawing this distinction is to distinguish between the reasons *there are* to φ and the reasons one *has* to φ.[49] Or, more fully, since normative reasons are always relational—reasons are always reasons *for* someone to do something, with the (possibly generic) subject being at least implicitly understood—the proposed distinction is between the reasons there are for a subject to φ and the reasons that subject has to φ.[50] Importantly, it's only the

---

[48] Cf. Broome (2007b: 352).

[49] This terminological distinction is also drawn by others (e.g. Williams (1986), Schroeder, etc.), though their understanding of it differs considerably from my own, rather schematic account.

[50] A similar distinction arises with talk of evidence—we can intelligibly talk of there being evidence for something (at some time *t*), even if that evidence hasn't yet been discovered (at *t*).

latter that are of direct relevance to rationality. To illustrate: suppose you're hungry and strongly desire pizza. As it turns out, there is one and only pizza parlor nearby: Moe's. Given this, it seems plausible that there's a reason for you to go to Moe's. However, from the fact that *there is* a reason for you to go to Moe's, it doesn't follow that you actually *have* any reason to go to Moe's. For suppose you have no idea that Moe's serves pizza, or that it serves food at all. In that case, even though, given all the facts, there's still something to be said in favor of you going there, from your perspective—given only what *you* have to go on—there isn't. In such a case there doesn't seem to be any sense in which your failing to (intend to) go to Moe's is irrational.

Something similar seems true in the epistemic case. For suppose you have skin cancer that has yet to be diagnosed. Unbeknownst to you, there are small spots forming in the middle of your back that are a highly reliable indicator of cancer. Although the fact that the small spots are forming in the middle of your back is a reason to believe you have skin cancer, it's not a reason you intuitively *have*—it's not something you're in a position to take into account in deciding what to believe, nor anything else that plausibly makes a difference concerning what it's reasonable for you to believe. So although your failing to believe you have skin cancer may be unfortunate, it wouldn't be in any meaningful sense *irrational* so long as you remain (non-culpably) ignorant of your condition.

As I said, the distinction between reasons there are and reasons had is a popular one. And it's intuitive, as the examples above illustrate. But it's also admittedly a bit facile. Even though it's a distinction we can mark *using* ordinary language (at least with appropriate emphasis and the right sort of examples), it's not one ordinary language itself insists upon—claims about there "being" a reason for someone to do something and that person "having" a reason to do something are oftentimes, and quite

unproblematically, treated equivalently.[51] Nonetheless, I do think there's considerable intuitive pull to the idea that reasons (like evidence) can't be relevant to rationality unless they're "had" in some sense—that is, unless they are somehow importantly related to, or part of, one's perspective.[52] So I do think there's an important distinction to be drawn, and it's one that is latent in (but not unambiguously marked by) our ordinary reasons-talk. It's a distinction between different kinds of normatively relevant considerations—or, more accurately, between different kinds of normative relations or statuses—one of which underlies our judgments about what there is reason to do or believe given some (largely)

---

[51] Indeed, there's a good case to be made that, linguistically, they *are* equivalent. In particular, it's plausible that in constructions of the form '*S* has (a) reason to φ' the subject '*S*' binds an implicit variable, the pronominal subject (which syntacticians call 'PRO') of the infinitival 'to φ', and that this variable is made explicit in the corresponding 'there is' constructions. Sæbø's (2009) analysis of 'have'-sentences then predicts that the following are equivalent:

(i) *S* has (a) reason to φ.

(ii) There is (a) reason for *S* to φ.

Of course, we don't say things like '*S* has (a) reason for her(self) to φ', but something similar holds for otherwise parallel cases like '$S_i$ has two sisters (of hers$_i$)', where the relational noun introduces an implicit/explicit variable (with the subscripted 'i' indicating coreference). (i) and (ii) would also be predicted to be equivalent to:

(iii) For *S*, there is (a) reason to φ.

However, as Broome (2013: 65) notes, there may be an additional reading of (ii) that is not equivalent to (i) and (iii). Consider:

(ii*) There is a reason for Carl to be punished.

This doesn't seem equivalent to:

(i*) Carl has a reason to be punished.

(iii*) For Carl, there is a reason to be punished.

According to Broome, (i*) and (iii*) ascribe "ownership" of the reason to Carl, whereas (ii*) does not. He thinks the two readings of (ii) can be distinguished grammatically:

> When 'There is a reason for *N* to *F*' ascribes ownership, the preposition 'for' governs '*N*'. When this sentence does not ascribe ownership, it governs '*N* to *F*'. The difference is invisible on the surface, but there is a test to detect it. When 'for' governs just '*N*', the phrase 'for *N*' can be shifted to a different place in the sentence without changing the sentence's meaning. When 'There is a reason for *N* to *F*' ascribes ownership, it means the same as 'For *N*, there is a reason to *F*'. When the sentence does not ascribe ownership, the shift is not possible. (65)

[52] The "subject-centeredness" or "perspectival" nature of rationality is widely attested; indeed, it may be the least controversial feature of rationality (which isn't to say it's *completely* uncontroversial). [Citations.]

perspective-independent body of information, and hence have nothing (directly) to do with what I'm calling reasons rationality, and the other of which underlies our judgments about what there is reason to do or believe given some (largely) perspective-dependent body of information, and hence has a direct bearing on reasons rationality.[53] Or really, since it's a distinction that cuts across the practical/theoretical divide, it's probably best understood as a distinction between different *groups* of normative relations or statuses.

To sum up: the distinction between 'reasons there are' and 'reasons had' is an intuitive, but ultimately technical one, introduced to mark a difference between two kinds of normatively relevant considerations, only one of which has a direct bearing on rationality. The important point is not semantic—one can mark the difference however one likes. But once we do, it's important to be sensitive to it in our theorizing. As will become clear below, there are a lot of complications concerning how the distinction should be properly understood, and not merely labelled.[54] For present purposes, however, I'll be ignoring such complexities, and rest content with the idea that the reasons relevant to reasons

---

[53] Another common way of marking (something like) the distinction between reasons there are and reasons had to distinguish between "objective" and "subjective" senses of 'reason', as well as the corresponding 'ought's they give rise to. As Alan Gibbard (2005) notes:

> We can ask what one ought to do in light of all the facts. Alternatively, we can ask what one ought to do in light of available information…Standardly in moral theory, we distinguish what a person ought to do in the *objective* sense and what she ought to do in the *subjective* sense. (179)

Although the objective/subjective distinction is some ways a natural one, I avoid it because it's misleading: facts about what reasons you have, or what you ought to do "in light of the available information", are just as "objective"—in the sense of being genuine, substantive, or real—as facts about what reasons there are, or what you ought to do "in light of all the facts". More than that, however, it's simply inadequate. And that's because the standard way (or, really, ways) of using 'objective' and 'subjective' is insensitive to the distinction between reasons rationality and structural rationality, and a further distinction would therefore be needed *within* the so-called 'subjective' in order to capture the difference between, say, the 'ought' of reasons rationality and the 'ought' of structural rationality.

[54] I'll return to this point in Chapters 6 and 7. There I argue that reasons rationality is not best understood in terms of reasons at all. But a closely related distinction is recoverable.

rationality—the reasons one "has"—are those things (facts, states, whatever) which play a certain characteristic role in determining what it is reasonable to do or believe. In short, they are what contribute to the reasonableness of actions and attitudes.[55]

The upshot is that once we become sensitive to some of the subtleties in our (normative) reasons-talk, we're able to diagnosis why counterexamples like Broome's are unconvincing, since they invariably involve appeal to reasons the existence and/or significance of which is outside the agent's ken.[56] The proposed counterexamples thus miss their mark because reasons rationality is only plausibly thought to be a matter of responding correctly to the reasons one has (understood per above), not all the reasons there in fact are. The fact that the fish in front of you contains salmonella certainly *is* a reason for you not to eat it, but so long as you remain wholly (and non-culpably) ignorant of this fact, it's not a reason you have—it's not something that intuitively makes a difference concerning what it's reasonable for you to do. It might be unfortunate if you ate the fish, but it wouldn't be irrational.

Broome (2007a: 167) himself recognizes more nuanced characterizations of the reasons relevant to rationality might be offered:

> Many philosophers find this quick objection convincing. As a result, few support precisely the idea that rationality consists in responding correctly to reasons. Instead they support a modified idea that connects rationality with the reasons you believe to exist rather than with actual reasons. They say that rationality consists in responding correctly to the reasons that you believe to exist, or something of that sort.

---

[55] Cf. Gibbons (2010).

[56] Broome doesn't offer a concrete counterexample in his (2013) book. Instead, he just provides a general description of the kind of cases that (he thinks) constitute counterexamples.

The alternative he considers, however, is importantly different from the alternative I (and others) favor.[57] For them to be equivalent, the following would have to be true:

(**HB**) You have a reason to φ iff you believe that a reason to φ exists.

Both directions are problematic. Consider first the left-to-right direction. I think it's pretty plausible that you can have a reason to φ without believing that you have one, or even that one exists. This is most obvious with young children and certain kinds of animals, both of which may have reasons to do (and/or believe) things but who may very well lack the concept of a reason and hence be incapable of *believing* they have reasons to do those things. There are also plausibly cases involving agents (like ourselves) who do or believe various things unreflectively, reflexively, or merely out of habit, without ever considering the question—even subconsciously—of what might be said in favor of doing or believing what they do. More fancifully, suppose I have a toothache and I'm told by a trusted, reliable friend that I should go to the dentist. But suppose I hold dentists in very low esteem, despite having received quality dental treatment in the past; I think (rather unjustifiably) that all dentists are incompetent, and that visiting one will only make matters worse. So I don't believe I have a reason to visit a dentist—I don't even think I have a *weak* reason to go, one that happens to be outweighed. On the contrary, I think the only reasons that bear on whether I should visit a dentist are ones that weigh *against* going. Nonetheless, given my toothache and my friend's testimony, I think I *do* in fact have a reason to go, and I'm not justified in not

[57] He actually considers another alternative as well—one which takes our attitudes themselves (our beliefs, intentions, and the like) to be the normative reasons rationality requires us to respond correctly to. He calls these 'attitudinal reasons', but is (rightly) dubious of their existence.

going. The left-to-right direction of (HB) therefore fails: one can have a reason to φ without believing that there is a reason for one to φ.

The right-to-left direction fails as well. Suppose that—contrary to all the evidence—I believe I'm the mayor of New York City, and so believe that there's a reason for me to go to City Hall. It doesn't follow that I thereby *have* a reason to go to City Hall, in the relevant sense; it's not the case that, given what I have to go on, anything actually counts in favor of going to City Hall. Doing so wouldn't be the least bit reasonable. It may of course may be *structurally* rational, given everything I believe; (mere) beliefs about reasons can certainly make a difference in what it's structurally rational to do or believe. But that's a separate matter. In the example at issue I merely *believe* that something counts in favor of going, and here as elsewhere, merely believing doesn't make it so.

To sum up the section so far: reasons rationality is concerned with normative reasons—i.e. things that count in favor of other things. But not with *all* normative reasons. For it's common to distinguish between two kinds of normative reasons, or normatively relevant considerations, only one of which has a bearing on what attitudes it is reasonable (and in that sense rational) to hold. These are the reasons one has (in the quasi-stipulative sense above), and their defining characteristic is that of contributing towards the reasonableness of—and in that sense "counting in favor of"—certain attitudes or courses of action.

## 1.8 Against unification

I've registered my conviction that reasons rationality and structural rationality are distinct, and that they come apart in both directions—one can be reasons rational without be structurally rational, just as one can be structurally rational without being reasons rational. Nonetheless, there have certainly been

attempts at bridging the divide. In epistemology, for instance, there are coherentist theories of justification, according to which for a belief to be epistemically supported or justified is for it to "cohere" with or be related in the right sort of way to one's other beliefs, or doxastic attitudes more generally.[58] And in the literature on practical reason, there are various "rationalist" attempts at explaining reasons for action in terms of (something like) structural rationality.[59] As Kolodny (2005) notes, such attempts "begin with an ideal of the rational agent, understood as one whose attitudes either stand in certain structural relations, or result from certain formal procedures…[a]nd they then understand reasons for action in terms of what such an agent would [or does] desire or will" (510).

The sort of unification view I'll focus on here is perhaps the most promising, because it tries to give a unified account of rationality across the board—both practical and theoretical—and in a way that incorporates features of *both* reasons rationality *and* structural rationality. The version I'll focus on is that developed by Mark Schroeder in a series of recent papers, though his view shares important similarities with, and is indebted to, that offered by Parfit (1997, 2011).

At the heart of Schroeder's account is the distinction between "objective" and "subjective" reasons, which is related but not equivalent to the distinction marked above between the reasons there are for an agent to $\varphi$ and the reasons that agent has to $\varphi$. For although the first half of each distinction—i.e. objective reasons/reasons there are—amount to the same thing, and hence line up nicely, the second half of each distinction—i.e. subjective reasons/reasons had—do not. For according to Schroeder, *S* has a subjective reason *r* to $\varphi$ just in case (i) *S* has a "presentational" state (i.e. belief or experience) with *r* as

---

[58] Some forms of "Subjective" Bayesianism can also be understood in this way.

[59] This is meant to include "Humean" accounts like that of Smith (1994) and "Kantian" accounts like that of Korsgaard (1996), though there are many (*many*) important differences between them.

content and, roughly, (ii) *r* would be an objective reason for *S* to *φ*, if *r* were true.[60] One way of understanding this proposal is as a substantive account of (what I'm calling) reasons had—that is, as a substantive account of the reasons that are relevant to reasons rationality. But the dialectic is tricky because Schroeder doesn't recognize the distinction between reasons rationality and structural rationality, and it's clear that he takes the notion of a subjective reason to be relevant not just to traditional debates in epistemology about knowledge and justification but *also* to the more recent debates about the nature and structure of coherence requirements.[61] That's why I think it's best, overall, to see him as a unifier—as someone who denies there's a deep difference between reasons rationality and structural rationality, and thinks a unified account of our intuitions concerning both can be provided.

Schroeder's notion of a subjective appears, at first blush, to perfectly suited to the task of unification, incorporating traditional elements from both reasons rationality and structural rationality. The appeal to *contents of belief* is clearly well-suited to account for many of our intuitions about coherence, and the inclusion of (contents of) *experiences* helps account for many of our intuitions about reasons. But there are problems. I'll consider five.[62]

First, Schroeder's account is restricted to presentational states and hence lacks the resources to explain the relevance of *non-presentational* doxastic states, such as agnosticism or suspension of judgment. This is especially important for capturing the (in)coherence data—an agent who believes *p*, believes that *p* entails *q*, but suspends judgment with respect to *q* is structurally irrational, and the suspension of judgment seems to be on all fours with the two beliefs in the explanation of why. Such an

---

[60] This is differs from Parfit's theory mainly in the kind of mental state that figures in clause (i)—Parfit restricts it to beliefs while Schroeder (2011) broadens it to include experiences as well. This avoids an important objection to Parfit's view (namely, that it doesn't account for the rational impact of experience), but it inherits the rest.

[61] [Citations]

[62] I'm indebted to Jim Pryor for several of these observations.

agent *will* be doing something wrong, according to Schroeder, since his beliefs furnish him with subjective reasons (namely, that *p* and that *p* entails *q*) that provide decisive support for believing *q*. But this seems to misdescribe what's going wrong with agent. As we saw above with the wholesale structural denier, it's not enough to provide an account of cases of incoherence that entails *something* has gone wrong; it's equally important that one provide an account that does justice to our sense of *what* has gone wrong, or else offer a plausible error theory as to why our judgments are mistaken.

The example involving suspension of judgment also illustrates how Schroeder's account fails to do justice to our intuitions concerning reasons rationality as well. For it's not at all obvious which attitude, if any, such an agent has most reason to give up—perhaps it's the belief that *p*, perhaps it's the belief that *p* entails *q*, or perhaps instead it's the suspension of judgment. Or perhaps it's all three. Merely specifying the relevant attitudes doesn't seem to settle the matter. Yet on Schroeder's view it does—as we've seen, the content of the agent's beliefs strongly (if not decisively) support believing *q*.

Just as Schroeder's account seems incapable of explanation the rational relevance of *non*-presentational doxastic states like agnosticism, so it seems incapable of explaining the relevance of *graded* presentational states such as degrees of confidence or experiences of varying quality. I might be fairly confident that *p* and that *p* entails *q*, for example, but not be confident enough to actually believe either and so not have them as subjective reasons to believe (or disbelieve) anything else. Yet clearly something would be wrong if I wasn't at all confident that *q*—and it would be even worse if I were believe that not-*q*. Since it's restricted to "all-or-nothing" presentational states, however, Schroeder's account can't make sense of any of this. A related problem concerns the justificatory impact of perceptual experiences. For it's plausible that the amount of justification experiences provide can vary depending on a number of factors, including quality, which can themselves vary. A higher-quality

experience as of *p*, for example, should provide more support for believing *p* than a lower-quality experience. Schroeder's account, in contrast, seems to treat all perceptual experiences on a par.[63]

A third worry concerns the rational relevance of *non-doxastic* (and non-presentational) states such as preferences and desires. The only way preferences and desires could furnish subjective reasons is by being believed, but this is odd, and many will insist that the first-order states themselves are capable of making a difference to what's (at least practically) rational—they needn't be filtered through one's beliefs to play such a role. (This is of course compatible with one's beliefs about one's preferences and desires *also* playing a role, at least when it comes to structural rationality.)

The fourth problem is that Schroeder's account is too content-focused: you can have a subjective reason to $\varphi$ only when you have a certain kind state with an appropriate content. But it's at least arguable that *non-representational* (not just non-*presentational*) states, such as headaches and other pains, are themselves reason-givers, and can thus make a difference concerning what it's rational to do or believe in the sense of reasons rationality. Such states don't have propositional content, however, and so on Schroeder's view they'd need to be believed to make a rational difference. Perhaps it's usually true that when we're in pain we also believe we're in pain, and so perhaps there's not a huge extensional difference being the first-order states themselves and our believing we're in those states; nonetheless, as before, many philosophers will insist that the first-order states *themselves* are capable of making a difference to what it's rational to do or believe, and hence that Schroeder mis-locates the source of rational support in such cases. In general, there seems to be an important difference between someone

[63] Perhaps Schroeder could insist that although quality of experience doesn't make a difference to whether or not it's content serves as a reason, it does make a difference to how *weighty* a reason it is. This may go some way towards addressing the worry about the variable impact of experiences, but the problem involving degrees of confidence remains untouched.

who actually has a certain kind of experience and someone who *merely believes* they're having such an experience, and this is a difference that the notion of a subjective reason is insensitive to.

The fifth and final worry is the most flat-footed, and that's that the notion of a subjective reason obscures the intuitively plausible distinction between (the contents of) mere attitudes and the reasons one has. It just seems wrong, for instance, to think that *merely* believing something puts you in a better epistemic position vis-a-vis its consequences that you would be otherwise, or that merely in virtue of *believing* that Obama is a Muslim and that all Muslims are terrorists one thereby has any *reason* (justification, evidence)—*pro tanto* or otherwise—to believe that Obama is a terrorist. Similarly, it seems implausible that merely *believing* that your child is the brightest student in the classroom gives you any reason or justification to be proud.[64] Such beliefs are clearly relevant to structural rationality, but not to reasons rationality. Having genuinely normative reasons, like having children, is harder than that.

So perhaps unsurprisingly the greatest theoretical strength of the notion of a subjective reason—namely, its potential in providing a unified account of rationality—also ends up being its greatest liability, and it is particularly ill-suited to explain the cases considered at the outset that motivate the need to distinguish between dimensions of rationality in the first place. Schroeder's account doesn't have the resources to do justice to *either* dimension of rationality, let alone both.

## 1.9 Plan for dissertation

Given the distinction between dimensions, the question now arises: how, exactly, are we to understand the nature of and difference between them? Although various answers are possible, the most prominent

[64] I should note that Schroeder (2011) addresses this objection, arguing that the "lack of justification guarantees defeat" and offers an error theory to explain away our intuitions that merely believing something doesn't, by itself, generate any justification (i.e. give us a reason) to believe it's consequences. Even if one found his defense convincing, though, the other problems would remain.

proposal locates a key difference between them in the structure of the requirements corresponding to each. For it is commonly thought that the domain of structural rationality consists of a distinctive set of "wide scope" requirements, while reasons rationality is thought of as either consisting of a set of "narrow scope" requirements or else as not fundamentally consisting of requirements at all.[65] For unlike requirements, the normative notions at the heart of reasons rationality (reason(s), evidence, justification, etc.) are not strict: you're not automatically irrational in virtue of not doing what you have reason to do—after all, you may have equal or more reason to do something else instead—whereas you *are* automatically irrational in virtue of violating a requirement of rationality.

As I'll argue, however, the dominant account of the difference between dimensions is mistaken. Indeed, it's mistaken twice over. For not only is it a mistake to think of the dimensions as differing with respect to the structure of their requirements, it's also a mistake to think of *either* domain as fundamentally consisting of requirements. We shouldn't understand structural rationality in terms of requirements any more than we should reasons rationality.

The plan for the dissertation is as follows. In Chapter 2, I introduce the debate over the nature and structure of rational requirements, with an emphasis on structural (or coherence) requirements. In Chapter 3 I undermine one of the main motivations for taking the scope debate seriously and in Chapter 4 I argue that there are no requirements of rationality in the relevant sense. The search for candidate requirements inevitably results in either falsehood or triviality. In Chapter 5, I provide an alternative account of the nature of and distinction between dimensions. Rather than strict requirements, I take each dimension to be characterized by a distinct and autonomous kind of pro tanto rational *pressure* or *force*—reasons rationality by (what I call) *justificatory pressure* and structural rationality by *attitudinal*

[65] Other than perhaps general ones, that is, such as those that implore you to do what you have most reason to do and believe in accordance with your evidence.

*pressure*. In Chapters 6 and 7 I the notion of a normative reason takes center stage, and I point out various subtleties in how we ordinarily think and talk about reasons—subtleties which, if taken seriously, have various upshots, both substantive and methodological.

# CHAPTER 2: RATIONALITY AND REQUIREMENTS

In this chapter I introduce, motivate, and explore the dominant requirements-based account of structural rationality. I begin by motivating the *existence* of rational requirements and then proceed to canvass some important foundational issues concerning their nature (what they are) and form (how they are structured). These further issues have generally been neglected and under-explored. This chapter is thus both expository and exploratory. The ultimate goal, of course, is to argue against the requirements-based picture, but it's important to first gain a clearer understanding of the target. I also have a more subversive goal: by taking the ideology of requirements more seriously than many of its proponents, I also hope to begin fostering suspicion of it. Again, ideology is not innocent.

## 2.1 The existence and explanatory status of rational requirements

The contemporary debate over structural rationality starts, innocuously enough, with the observation that which attitudes we actually have—whether or not they're justified (or "good")—can make a difference concerning what other attitudes it's structurally rational for us to have. For example, as we've seen, what you believe can make a difference concerning what else it's rational for you to believe, and what you intend or desire can make a difference concerning what it's structurally rational to do. (For ease of expression, I'll often drop the 'structural(ly)' qualifier in what follows, though for clarity's sake it will occasionally re-appear.) More than merely making a difference, however, it seems that having certain attitudes can *rationally commit* you to having other attitudes, in such a way that you will definitely do something wrong or exhibit a rational failing if you have the former without having the latter. Similarly, it seems that having certain attitudes can *rationally prohibit* you from having other attitudes, in such a way

that you will definitely do something wrong or exhibit a rational failing if you have the former attitudes while also having the latter. For this reason, many philosophers find it natural to think of the domain of structural rationality as consisting—at least in part—of a distinctive set of *rules* or *requirements* that mandate or prohibit certain combinations of attitudes, in much the same way that the law consists of a distinctive set of legal rules or requirements that mandate or prohibit certain actions.[66]

The most prominent and influential proponent of a requirements-based conception of structural rationality is John Broome (1999, 2004, 2007, 2013), but he's just one of many who take talk of rational requirements seriously. The basic idea is that for each kind of incoherent combination of attitudes there is a corresponding rule or principle prohibiting it, and that what's wrong with incoherent agents is that they violate these principles—they're irrational in virtue of violating these principles—just as for each illegal action there is a law that prohibits it, and in virtue of which actions of that kind are illegal. Suppose that Jill steals Jack's bike. What Jill did was illegal. But why exactly? Subtleties aside, the answer is clear: Jill took Jack's bike without his permission, and there's a law that prohibits taking someone else's property without their permission.[67] If there hadn't been a law prohibiting theft, then although what she did may have been immoral, it wouldn't have been illegal.

The requirements of rationality are standardly thought to play an analogous role: just as laws explain why particular (kinds of) actions are legal or illegal, so rational requirements are supposed to explain why particular (combinations of) attitudes are rational or irrational. Tom, above, is irrational.

---

[66] Morality is sometimes thought to be constituted by certain rules or principles as well, though I'm just as skeptical of a requirements-based conception of morality as I am of rationality.

[67] This is how explanations of particular normative facts work in general—they typically involve appealing to particular non-normative facts together with (more) general normative facts, though the nature of the general normative facts and what, if anything, explains them may vary from case-to-case or from one domain (e.g. legal, moral, epistemic) to another.

But why exactly? According to the present line of thought, it's because he fails to believe the obvious consequences of other things he believes, and there's a requirement that prohibits him (and us) from doing so.

The requirements of rationality are thus supposed to be *more* than mere necessary conditions for being fully rational, and in two ways. First, it's not enough that a subject fails to be rational whenever the requirements are not met; the requirements, when violated, are supposed to guarantee something more—namely, irrationality. And we obviously shouldn't conflate being irrational with merely *not* being (fully) rational. Rocks and trees fail to be rational, after all, but they're not thereby *ir*rational. They're *a*-rational—they lack the relevant sort of complex capacities needed in order for them or their states to be apt candidates for *any* sort of rational evaluation. In general, if *C* is a necessary condition for being rational, then although it follows that not-*C* is a sufficient condition for *not* being rational, it doesn't follow that not-*C* is a sufficient condition for being *ir*rational. Unlike mere necessary conditions, rational requirements are supposed to state conditions such that "violating" them guarantees the having of a negative evaluative property—namely, being irrational—and not just the lack of a positive property—namely, being rational. But rational requirements are also supposed to do more than merely *guarantee* irrationality. In particular, they're supposed to state conditions such that if an agent fails to meet the conditions, they are *thereby* irrational—the agent is irrational (at least partly) *in virtue of* violating the requirement. The requirements are thus supposed to be revealing something about the *nature* of structural rationality, rather than merely providing a diagnostic for it. The general point should be a familiar one. For in providing an explanatory account of something—whether it be why WWII occurred, what it is to be conscious, the way gravity works, or the nature of mathematical facts—we're almost

always interested in providing something *more* than mere necessary conditions. The same goes for providing an account of structural rationality.

Different authors express the point in different ways. For instance, Jonathan Way (forthcoming) 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.

In a similar vein, Broome (2013) distinguishes between two senses of 'requires' and its cognates: a weaker "property" sense and a stronger "source" sense.[68] He thinks the most interesting questions concern what rationality requires in the latter sense, not the former. (We'll see why shortly.) Schroeder (2013) agrees, and takes his own talk of the "rules" of rationality to be equivalent to Broome's notion of a source requirement:

> [Y]ou count as (having the property of being) irrational *in virtue of* breaking one or more of the rules (source requirements) of rationality, and [the debate concerns] which rules (source requirements) you are breaking…when you have inconsistent beliefs, are akratic or means-end incoherent, or fail to draw the obvious consequences from your other beliefs. (299, emphasis added)

---

[68] As he puts it:

> The first appears in constructions where its subject denotes a property: 'Beauty requires hard work'; 'Staying healthy requires you to eat olives'; 'Success in battle requires good horses'; 'Crossing the Rubicon required determination'…[The second appears in] constructions [where] the subject of 'requires' denotes a person or thing that has some sort of real or presumed authority: 'The minister requires the ambassador's presence'; 'The law requires you to drive carefully'; 'The bill requires payment'; 'Fashion requires knee-length skirts'; 'My conscience requires me to turn you in'. (109)

I'm belaboring this point about the explanatory status of the requirements (or "rules") of structural rationality because *it's* the claim I ultimately want to reject. It would be foolhardy to deny there are necessary conditions of being fully rational, and in *that* sense some "requirements" of rationality. After all, in order to be rational one has to be alive, but being alive isn't rationally required in the relevant sense.[69] Instead, what I'm concerned to deny is that there are requirements of rationality in the *strong* sense. The nature of structural rationality—just like the nature of reasons rationality—is importantly *un*like that of the law.

Despite the emerging consensus concerning the status of structural requirements as explanatory, however, there is consensus about little else. One important but neglected question concerns the "source", or explanation, of the laws and requirements themselves. To characterize something as a law or requirement is to attribute to it a certain normative status, and this raises the question of what it is (if anything) that explains or is responsible for its having that status: how do they arise, and from whence do they derive their authority? Call this the "source question". The answer will presumably vary across domains. In particular, while it's fairly clear (at least in broad outline) what makes a law a law, it's far less clear what the source of rational requirements might be.[70] Broome (2013) considers this question but doesn't provide an answer; Schroeder (2013) likewise considers it, and argues (a lá Kant) that

---

[69] Cf. Broome (2007, 2013). Nor is it rationally required that 2 + 2 = 4, even though that too (along with any tautology or other necessary truth) is, trivially, a necessary condition of being rational—it's impossible for one to be irrational without it being the case that 2 + 2 = 4. So it should be uncontroversial that not all necessary conditions are created equal. Of course there *are* non-trivial questions concerning the conditions that have to be met in order to be rational, and some are only concerned with what rationality "requires" in this weak sense. See, for example, Titlebaum (2013, 2015) and Fitelson and Easwaran (2015). But for each non-trivial necessary condition proposed, there's a further question to be answered—namely, *why* is it a necessary condition? What explains its (non-trivial) status as necessary? A complete story needs an answer.

[70] [Citations concerning nature of the law, legal positivism, etc.]

they're rules we impose on ourselves. I reject a requirements-based conception of rationality across the board, so I don't face the source question, though I do face a similar question concerning the "source" of rational pressure. I gesture at the beginnings of an answer at the end of the Chapter 5.

There are many other important questions that proponents of a requirements-based picture of structural rationality might—and do—disagree about. Besides the question of their status and source, these questions include ones their *content*, *scope*, *jurisdiction*, *extent*, *temporal nature*, and *normative status*. Few of these have received the attention they deserve, and it's common for different questions to be conflated. Much of this neglect is due to the fact that the ideology of requirements is usually assumed to be common ground between supporters and detractors alike, and this has discouraged critical evaluation of the baggage the ideology carries with it. I'll consider each in turn.

## 2.2 The content of requirements

First, there's the question of **content**. That is, for each putative coherence requirement *R*, there's the question: what exactly is *R*'s content? For example, in order to be deductively coherent, does your set of beliefs have to be closed under logical entailment? That is, do you have to believe all of the logical consequences of one's beliefs, even those that are far from obvious or those outside your ken altogether? Many think not, including those who are otherwise sympathetic with the thought that there's a close connection between logic and rationality. Even if believing all the logical consequences of one's beliefs is in some sense a rational ideal—and it's far from obvious that it is[71]—it's commonly thought that it's simply too demanding to be understood as a rational *requirement* for cognitively limited

[71] For a particularly careful and compelling objection to deductive closure as a rational ideal (not to mention requirement), see Worsnip (forthcoming-c). This is compatible, of course, with deductive closure being useful as an *idealization* for certain theoretical purposes. No one should object to that.

creatures like ourselves. In general, failing to satisfy a standard that is not humanly possible to satisfy, even under optional conditions and with the benefit of increased (but still recognizably human) cognitive powers, is not a mark of irrationality in any genuinely normative sense.[72] As a result, it's standard to restrict the relevant requirement in such a way as to only concern logical consequences that are sufficiently "obvious" and/or "relevant".[73]

Though important, and required to maintain plausibility, such an amendment raises a number of vexed issues of its own. For example, do the consequences need to be *actually* obvious, or is enough that they be *potentially* obvious—that is, such that they *would* be obvious upon casual (or at least careful) reflection? And *to whom* do they need to be obvious—the agent herself as she's currently constituted, or instead some suitably idealized version of the agent, or what? I doubt there's a general answer to these questions, but even so there's a more fundamental problem. And that's that obviousness *comes in degrees*. As a result, any requirement formulated in terms of some (possibly vague) threshold of obviousness will inevitably seem arbitrary, and won't (on its own) accurately reflect the gradable nature of the underlying phenomenon.

This sort of worry about the inability of "threshold-y" (or all-or-nothing) requirements to do justice to the underlying phenomenon will take center stage in Chapter 4, where I'll argue against the existence of rational requirements. But for now I'll leave it to the side. And that's because there's another important question to be addressed in connection with the proper understanding of the relationship between logic and (structural) rationality—namely, what about *mere beliefs* about logical relations such as entailment? Do they make a rational difference? It's natural to interpret claims about

---

[72] This is one application of Schechter (2013) calls ''the Spiderman principle'' in epistemology: With greater cognitive power comes greater epistemic responsibility.

[73] See, for example, the exchange between Field (2009) and Harman (2009).

"obviousness" as factive—that is to say, if *q* is an obvious consequence of *p*, then it seems to follow that *q* *really is* a consequence of *p*.[74] But what if it merely *seems* obvious to you that *q* is a consequence of *p*, even though as a matter of fact it's not a consequence? Or what if you *merely believe*—perhaps on the basis of testimony—that *q* is a consequence of *p*, even though it's not? Clearly you can be wrong about the logical relations between propositions, including those you believe. And if you are, won't such beliefs make a difference concerning what else it's structurally rational for you to believe, just as your other (non-logical) beliefs do? It's not clear what grounds we would have for treating beliefs about logical entailment differently than other beliefs, and I, for one, find it hard to countenance the idea that (possibly false) logical beliefs make no rational difference whatsoever. So I'll be proceeding on the assumption that mere beliefs of all stripes—whether or not they're true and whether or not they're justified—are relevant to structural rationality.

This is not to deny that other things besides beliefs (and other doxastic states) might *also* be relevant, including "mere" truths about logical relations. But it's worth comparing the analogous debate in the practical realm concerning means-end coherence. Suppose you intend to buy a bottle of wine for dinner. And suppose that, unbeknownst to you, the local liquor store just closed and that the only other place to buy wine is on the other side of town. Since news of the local liquor store's closure hasn't reached you, you assume it's still open, and so intend to buy the wine there on your way home from work. This seems like a perfectly sensible plan. Indeed, although the world has conspired against you, there doesn't seem to be any sense in which it you're rationally criticizable for having such a combination of attitudes. And this is true even though you fail to intend the means that, as a matter fact, are necessary to your ends. What examples like this strongly suggest, then, is that *mere facts* about which

[74] There may also be a use of 'obvious' that is more purely phenomenological, and not factive; I don't intend to be taking a stand on the question.

means are necessary (and/or optimal) to which ends are rationally irrelevant. As a result it's standardly assumed that all that matters, as far as rationality is concerned, is what agents *believe* about the means to their ends (or what their credences about such are), and/or what it would be *reasonable* for them to believe about the means to their end. Both of these can, and oftentimes do, come apart from the facts themselves.

It's very tempting to say the same thing when it comes to "mere" logical facts.[75] You might believe *p*, for example, and believe *q*, and yet in no way be irrational in failing to believe *r*, even though—as a matter of fact—*r* is a logical consequence of *p* and *q*, perhaps together with certain other things you believe. The "proof" might be incredible complex, and you might be in no position to construct it. Is there really any legitimate sense in which you're rationally criticizable, or incoherent, in failing to believe *r*? I doubt it. Again, this is compatible with saying that in failing to believe *r* you're falling short of some rational ideal or other. But falling short of some ideal doesn't, on it's own, merit criticism, or reflect poorly on you. In contrast, violating a genuine requirement of structural rationality *is* suppose to merit criticism—it *is* supposed to give rise to the sense that something has gone wrong with you (namely, you have incoherent attitudes), and it *is* supposed to reflect poorly on you in your capacity as a rational agent.

Despite my evident sympathies, however, I needn't take a definite stand one way or the other on whether "mere" logical facts and relations themselves make a difference to structural rationality. What's more important for my purposes is that *beliefs about* logical facts and relations can, and usually do, make a difference. If you believe *p* and believe that *p* entails *q*, for instance, and yet you're not sure about *q*, intuitively something has gone wrong—your attitudes fail to fit together properly. And this judgment

---

[75] See Harman (1986) for the classic defense of this view.

persists even if we suppose that, as a matter of fact, *p* does *not* entail *q*. Just as beliefs about the causal or constitutive relations between means and ends matter for means-end coherence, so beliefs about logical relations between propositions matter for deductive coherence.

As the above discussion illustrates, for any given requirement, the question of content is highly non-trivial. For the most part, however, my concern won't be with getting the details right. Instead, I'll be working with simplified versions of various (purported) requirements.

## 2.3 The scope of requirements

Second, there's the question of *scope*. That is, for each putative requirement *R*, there's the question: does *R* mandate or prohibit a particular combination of attitudes, or does it instead mandate or prohibit a specific attitude (at least under certain conditions)? For example, if you intend an end *E* and believe that in order to achieve *E* you have to take means *M*, does it follow that you are rationally required to intend *M*? Or are you merely required, at all times and irrespective of your other attitudes, to not have the following combinations of attitudes: intending *E*, believing that *M* is a necessary means to *E*, and not intending *M*? Similarly, if you believe you ought to φ, does it follow that you are rationally required to intend to φ? Or are you merely required, at all times and irrespective of your other attitudes, to not have the following combinations of attitudes: believing that you ought to φ and not intending to φ? In each case, so-called "narrow-scopers" think that the first claim is the intuitively correct verdict, whereas "wide-scopers" opt for the second claim.

There are controversies concerning the logical form of the corresponding requirements, as well as their exact content, but perhaps the most common and least controversial way of representing the

competing interpretations of (what I'll call) the requirement of Mean-End Coherence (ME) requirement and the Anti-Akrasia (AA) requirement are as follows, where '—>' denotes the material conditional:

(**ME-Narrow**) (You intend end $E$ ∧ you believe that $M$ is a necessary means to $E$) —> rationality requires of you that you intend $M$.

(**ME-Wide**) Rationality requires of you that ((you intend end $E$ ∧ believe that $M$ is a necessary means to $E$) —> you intend $M$).

(**AA-Narrow**) (You believe you ought to φ) —> rationality requires of you that you intend to φ.

(**AA-Wide**) Rationality requires of you that (you believe you ought to φ —> you intend to φ).

Since in (ME-Wide) and (AA-Wide) the (intentionally artificial) 'rationality requires' operator governs a material conditional, they can be re-expressed more naturally as the following:

(**ME-Wide'**) Rationality requires of you that you not: intend end $E$, believe that $M$ is a necessary means to $E$, and not intend $M$.

(**AA-Wide'**) Rationality requires of you that you not: believe you ought to φ and not intend to φ.

These do a better job at transparently capturing the core thought animating many wide-scopers—namely, that coherence requirements *ban* incoherent combinations of attitudes, and that's it. As a result, the requirements are simply silent as to which particular attitudes you should and shouldn't have.

With coherence requirements containing more than two conditions, there are also "intermediate scope" interpretations available. For example, rather than (ME-Narrow) and (ME-Wide), Way (2010) opts for the following:

(**ME-Intermediate**) You believe that *M* is a necessary means to end *E* —> rationality requires of you that (you intend *E* —> you not intend *M*).

Notice that although both (ME-Narrow), (ME-Intermediate), and (ME-Wide) involve the same three conditions, and in that sense the same content, they differ in how the conditions are related and, more specifically, whether the 'rationality requires of you' operator governs a particular attitude (as in (ME-Narrow)) or instead a combination of attitudes (as in (ME-Wide) and (ME-Intermediate)). The question of scope thus concerns the *form* or *structure* of requirements.

We face the same question concerning scope when we (re-)consider the question of what deductive coherence requires. I indicated above that there's debate about it's exact content, just as there is for the means-end requirement, and in particular whether mere logical facts and relations can make a rational difference. Even if they do, however, it seems that our *beliefs* about them can also make a difference, along with any other doxastic attitudes and/or dispositions we may have. So I'll focus on

fairly simple and widely accepted principles concerning the proper relationship between certain beliefs.[76] Here are narrow-scope and wide-scope versions of the requirement of (what I'll call) Deductive Coherence (DC):

(**DC-Narrow**) (You believe $p$ ∧ you believe that $p$ entails $q$) —> rationality requires of you that you believe $q$.

(**DC-Wide**) Rationality requires of you that ((you believe $p$ ∧ you believe that $p$ entails $q$) —> you believe $q$).

Given these formulations, we can now better appreciate a couple of reasons why it is often thought to be important that the requirements of rationality represent something *more* than mere necessary conditions. For suppose we interpret 'rationality requires of you that' in both of the above as meaning 'a necessary condition of your being (fully) rational is that'. In that case, (DC-Wide) will actually be *entailed* by (DC-Narrow)—any way of making the latter true will also be a way of making the former true. One way of making (DC-Narrow) true is to not satisfy the antecedent—that is, to either not believe $p$ or else not believe that $p$ entails $q$. But in that case (DC-Wide) will also be true—because the antecedent of the embeded material conditional will be false, the conditional as a whole will be true, and so you'll have satisfied the condition that according to (DC-Wide) it is necessary to meet in order for you to be rational. But now suppose you *do* believe $p$ and that $p$ entails $q$. Then according to (DC-Narrow)—on this interpretation—it is a necessary condition of your being rational that you believe $q$. But if it's a necessary

[76] Considering the range of other doxastic attitudes and the proper relationship between them is obviously also very important; but it's also much more complicated, and inessential for my purposes, so I'll continue working mostly with coarse-grained attitudes.

condition of your being rational that you believe *q*, then (trivially) it's a necessary condition of your being rational that you either believe *q or* not believe both that *p* and that *p* entails *q*—and hence that you satisfy the material conditional in (DC-Wide). After all, if it's a necessary condition of A that B, then it is a necessary condition of A that either B or C, for any C.

For similar reasons, (ME-Narrow) will entail (ME-Wide). So if we interpret the above requirements as simply stating necessary conditions, then (ME-Wide) and (DC-Wide) won't be controversial, but instead descriptions of the data that everyone is trying to explain.[77]

Even if we do interpret 'rationality requires of you that' as meaning 'a necessary condition of your being rational is that', there will still be a debate about whether, for example, (ME-Narrow) and (DC-Narrow) are *also* true, in addition to their Wide counterparts. However, as Schroeder (2014) notes with respect to (ME-Wide):

> [P]rominent objectors to the Wide scope view in the literature—including Niko Kolodny, Joseph Raz, and myself—have not presented their view as granting that [(ME-Wide)] is true but insisting that one of the other theses is also true. On the contrary, many remarks of these critics suggest that they have been trying to argue directly against [(ME-Wide)]. So it seems that charity requires taking them to interpret ['rationality requires of you that'] as meaning something else. (220)

The real debate, then, is over what *explains why* intending end *E,* believing that *M* is a necessary means to *E,* and not intending *M* guarantees irrationality. Is it because (ME-Wide) itself specifies a ground-floor requirement of rationality, as wide-scopers insist, or is it because by having the attitudes in the antecedent of (ME-Narrow) one is thereby required to have the attitude specified in the consequent, as

[77] Cf. Schroeder (2014: 220).

narrow-scopers insist?[78] And of course the point generalizes. This serves to reinforce the point above concerning the explanatory status of the requirements of rationality—without such an assumption, the history of the scope debate makes little sense.

Another reason the scope debate only makes sense if we take requirements to be something more than mere necessary conditions is that wide scope and narrow scope requirements make *exactly the same predictions* concerning when an agent has the property of being (fully, structurally) rational.[79] So if all we cared about where the conditions that need to be met in order for an agent to be rational, the scope debate would be explanatorily idle. But the property of being rational is *not* the only thing that matters. For one thing, as we've seen, we're also interested in the property of being structurally *irrational*, and this is not the same thing as the property of not being structurally rational—being incoherent is not just a matter of not being coherent. It's the violation of requirements that are supposed to explain why an agent has the property of being irrational, when they do. And although wide scope and narrow scope requirements make the same predictions concerning when an agent is irrational, they don't offer the same explanations of *why*—is the agent irrational because they have a combination of attitudes rationality requires them not to have, or instead because they fail to have a particular attitude that, in their present circumstances, rationality requires them to have? What's more, as we'll see in a moment, although wide scope and narrow scope requirements align when it comes to violation-conditions—that

---

[78] Notice that I said '*thereby* required' (emphasis added this time around). And notice that this bit of explanatory information is completely left out of (ME-Narrow), formulated as it is using the material conditional. This one of several ways in which (ME-Narrow) is not inadequate as a representation of the relevant narrow scope requirement. Indeed, I'd go further and say that it's positively *distorting* as a representation. We'll return to this issue later.

[79] See Broome (2007a, 2013) and Schroeder (forthcoming) for more on this worry.

is, the conditions under which they're violated—they differ when it comes to compliance-conditions—that is, the conditions under which they're complied with. Finally, as Broome (2013: 119) notes,

> We are [primarily] interested in what [rationality] requires of you in the particular situation you are in…You may be very far from having the [global] property [of being fully rational]…In practice our interests are local. We are more interested in coping with our actual imperfect situation than in how to be perfect.

I'll return to this point about the 'local' character of the requirements later.

## 2.4 The jurisdiction of requirements

This brings us to the third question a theory of requirements faces. And that's the question of *jurisdiction*, or domain of governance. That is, for each putative requirement *R*, there's the question: does *R* govern all rational agents at all times, or does it only apply under more selective conditions? This question is analogous to the one that arises in the legal realm. Schroeder (2014) puts the point as follows:

> One important feature of laws is that they have *jurisdictions*. For example, in the state of New York, it is illegal to turn right at a red light. The jurisdiction of that law is *drivers in New York*, and what it prohibits is *turning right on red*. In general, anyone who is simultaneously a driver in New York and is turning right on red is in violation of this law, but it is important to appreciate that being a driver in New York and turning right on red make different contributions to this fact. If you are a driver in New York and you *don't* turn right on red, then you are *complying* with the law, whereas if you are a pedestrian in New York or a driver in Buenos Aires or Cairo, the law simply doesn't apply to you. The reason why drivers in Cairo who turn right on red aren't in violation of New York traffic laws is that the New York state legislature doesn't have *jurisdiction* over drivers in Cairo—not that it does have jurisdiction, but that they are in compliance. (217)

We might actually want to draw further distinctions here. In particular, there seems to be a sense in which residents (or visitors) of New York *are* governed by, or subject to, the relevant traffic laws, even

when they aren't driving, whereas residents of Cairo are not. Those residing in New York are *ipso facto* subject to its laws, and in that sense fall within their jurisdiction. Those residing in Cairo, by contrast, are outside of its jurisdiction altogether, and so not subject to New York traffic laws in any meaningful sense. When it comes to New York state traffic laws there thus seems to be a categorical difference between residents of New York, whether or not they're driving, and residents of Cairo. Of course that's *not* to say that if you're a pedestrian in New York you bear exactly the same relationship to the traffic law as those who are driving do. For there's a clear sense in which the law isn't concerned with those who aren't driving, and so doesn't apply to you in your present circumstances as a pedestrian. This is nonetheless compatible with the thought that everyone residing in New York is prohibited from turning right at a red light while driving, and hence everyone in New York is within the jurisdiction of the relevant law, whereas the law is simply silent about those residing elsewhere.

To capture this difference, I think it's worth distinguishing between the *jurisdiction* of a given law —i.e. those which are "subject to" the law or within its domain—and the *conditions of application* of the law—i.e. those conditions that need to be satisfied by those within its domain in order for it to actually apply to a particular case. Violation and compliance require that the conditions of application obtain; merely being within the jurisdiction of a law and not satisfying it's conditions of application is not enough for compliance. Nor, of course, is falling outside of the jurisdiction of the law altogether. We should thus distinguish between (at least) three, rather than two, types of non-violation. A given person can fail to violate a given law *L* by:

(a) complying with *L* (e.g. not turning right at a red light while driving in New York)

(b) not satisfying *L's* conditions of application (e.g. not driving while in New York), or

(c) being outside of *L*'s jurisdiction altogether (e.g. being in Cairo).

This is relevant because, as Broome and Schroeder (and others) stress, the relevant concept of a rational requirement is supposed to be law-like not just with respect to its explanatory status but also insofar as it allows for a non-vacuous distinction between those within its jurisdiction and those who are not. This is what allows us to distinguish between two forms of non-violation—compliance and avoidance—which in turn is thought to allow for an intuitive "test" of whether a given principle should be interpreted as wide or narrow or intermediate scope. The dialectic is muddied, however, by a failure to distinguish between jurisdiction and conditions of application, and the concomitant failure to distinguish two forms of non-application—namely (b) and (c) above.

For example, Schroeder asks us to consider the following regimentations of the relevant New York state traffic law (which, like the requirements of rationality, are supposed to be stating more than mere necessary conditions for not doing something illegal):

| | |
|---|---|
| (**NY-Narrow**) | (You are in New York ∧ you are driving) —> New York state law requires of you that you not turn right on a red light. |
| (**NY-Intermediate**) | (You are in New York) —> New York state law requires of you that (you are driving —> you not turn right on a red light). |
| (**NY-Wide**) | New York state law requires of you that ((you are in New York ∧ you are driving) —> you do not turn right on a red light). |

He interprets (NY-Narrow) as claiming that the relevant New York state traffic law has jurisdiction over drivers in New York and requires of them that they not turn right on red. (NY-Intermediate) is then be interpreted as claiming that the relevant New York state traffic law has jurisdiction over everyone in New York, whether or not they're driving, and requires of them that they either not drive or else not turn right on red. And (NY-Wide) is interpreted as claiming that the relevant New York state traffic has jurisdiction over everyone, whether or not they're in New York, and requires of them that they either not be in New York, not be driving, or else not turn right on red.

On this interpretation (NY-Wide) is quite implausible, whereas (NY-Narrow) and (NY-Intermediate) fare better. Schroeder thinks this can be brought out by considering the conditions under which drivers count as complying with the relevant law as opposed to "avoiding" or "escaping" it:

> Another way of seeing the same thing [is] to observe that drivers in Cairo and Buenos Aires are not *complying* with New York state traffic regulations, simply because they are not in New York. In contrast, drivers in New York who don't turn right on red *are* complying with New York state traffic regulations. It is true that there are two ways to *avoid violating* New York traffic regulations —you can refrain from turning right on red, or you can leave the state. But these are not two ways of *complying* with the regulations. One is compliance, and the other is escape. The distinction between compliance and escape tracks the regulations' *jurisdiction*, because you can comply with a regulation only if you fall under its jurisdiction, and leaving the jurisdiction of the regulation is sufficient to avoid violating it.

Schroeder is right that you can avoid violating the New York traffic law by refraining from turning right on red or by leaving the state, and that only the former counts as complying with the law. But as noted above, there's also a third way of avoiding violating the law—namely, by being in the state but not driving. You don't intuitively comply with the relevant law merely by not driving while in being New York, but neither does one seem to be "avoiding" or "escaping" the law in the same way that you would

be if you were to leave the state altogether. Once we draw the further distinction between jurisdiction and conditions of application, we're able to offer a more natural characterization of the relevant law. And that's as the law having jurisdiction over everyone in New York but only applying to those within its domain who are driving. Those to whom it applies are then required not to turn right on red. This way of understanding the relevant law is akin to the above interpretation of (NY-Intermediate) insofar as its jurisdiction includes everyone in New York, and akin to the interpretation of (NY-Narrow) in its differential treatment of driving (a condition of application) and not turning on red (what's required). But both (NY-Intermediate) and (NY-Narrow) themselves are inadequate, since they don't discriminate between jurisdiction and conditions of application and so aren't able to distinguish between the different roles played by each condition.

So although Schroeder is right to emphasize that you can comply with a regulation only if you fall under its jurisdiction, and that leaving the jurisdiction of the regulation is sufficient to avoid violating it, he's wrong to think the distinction between compliance and "avoidance" tracks the notion of jurisdiction. The notion of jurisdiction allows us to distinguish those within the domain of the law—those who are "governed" by or "subject" to it—form those who are not—those who "avoid" or "escape" its reach. And the notion of a condition of application allows us to draw a further distinction within the law's domain between those to whom the law applies—and hence who complies with or violates the law—and those to whom it does not.

The upshot of all of this is that we're now in a position to better understand the nature of (purported) narrow scope requirements. Recall again the Narrow and Wide interpretations of the means-end (or instrumental) requirement:

(**ME-Narrow**): (You intend end $E$ ∧ you believe that $M$ is a necessary means to $E$) —> rationality requires of you that you intend $M$.

(**ME-Wide**): Rationality requires of you that ((you intend end $E$ ∧ believe that $M$ is a necessary means to $E$) —> you intend $M$).

Like Broome, Schroeder interprets (ME-Wide) as having universal jurisdiction over all rational agents, but unlike Broome he finds this puzzling:

> I have to confess that I find this idea as puzzling as I do the idea that the New York state legislature has jurisdiction over drivers in Buenos Aires. What is the source of this universal requirement supposed to be, and how does it acquire jurisdiction over every rational agent? I find it difficult to even get my head around this question.

This is a good question—it's the "source" question considered above, and it's one Broome (2013) acknowledges but doesn't answer. He thinks it's plausible to view the requirements of rationality as being on a jurisdictional par with those of morality, governing agents in general, rather than something like a state or federal law, governing only a particular subset of agents. Schroeder disagrees, and thinks we should instead accept either (ME-Narrow) or (ME-Intermediate), interpreted in such a way that their jurisdiction is limited to agents who satisfy their antecedent conditions. This, he thinks, allows for a more promising (neo-)Kantian interpretation according to which they specify requirements that we impose on ourselves. How does this help answer the source question? Here's Schroeder again:

> Whereas it is puzzling [how] something could come to have jurisdiction over every rational agent, I don't think it is similarly puzzling how a rational agent could come to have jurisdiction over herself. This is not to say that there are *no* philosophical puzzles about the latter…but only that the same puzzles do not arise. If an agent has jurisdiction over herself, then she can create rules or laws for herself…This picture does away with the idea that there is some peculiar source

> of rational requirements which somehow has jurisdiction over every rational agent, and replaces it with the idea that each rational agent has jurisdiction over herself. (224-225)

This looks to me like a clear case of swapping one mystery for another with little or no explanatory gain; at the very least, it's hard to see how it represents a genuine dialectical advance. So I won't pursue Schroeder's proposal any further here.

What's important to note, however, is that there's another—and to my mind more plausible—way of understanding narrow and intermediate scope requirements like those specified by (ME-Narrow) and (ME-Intermediate). It's a view according to which they, like (ME-Wide), have jurisdiction over all rational agents, but only *apply* to agents who satisfy their the conditions (i.e. have the attitudes) specified in their antecedents. That is, their antecedents specify conditions of application rather than delimit their jurisdiction. These views obviously face the same source question that wide scopers like Broome do, but they enjoy a couple important advantages over a more traditional understanding of the narrow scope view. I'll focus on (ME-Narrow), but similar remarks apply to (ME-Intermediate), as well as all other non-wide scope requirements interpreted in the way I'm suggesting.

The first advantage arises from the fact narrow scope requirements like (ME-Narrow) make different—and intuitively more plausible—predictions than wide scope ones like (ME-Wide) concerning what counts as compliance.[80] Consider a particular instance of (ME-Wide) and compare it to the corresponding instance of (ME-Narrow):

(**Grandma-Wide**) Rationality requires of you that ((you intend to visit your

[80] Lord (2014) calls this the "real symmetry problem" for the wide scope view, though it's importantly different from the original "symmetry problem" that Kolodny, Schroeder, Way, and others, have raised, a version of which I'll consider below. So they are really best thought of as two separate objections.

grandmother ∧ believe that driving to her house is a necessary means to visiting her) —> you intend to drive to her house).

(**Grandma-Narrow**) (You intend to visit your grandmother ∧ believe that driving to her house is a necessary means to visiting her) —> rationality requires of you that you intend to drive to her house.

Both (Grandma-Wide) and (Grandma-Narrow) have jurisdiction over all rational agents, but unlike (Grandma-Narrow), (Grandma-Wide) also *applies* to all rational agents. And what it requires is that you satisfy the embedded material conditional. Any way of doing so is thus a way of complying with the requirement. If you don't intend to visit your grandmother, you thereby comply with (Grandma-Wide), regardless of what else you believe or intend. And the same goes if you don't believe that driving to her house is a necessary means to visiting her, or if you happen to intend to drive to her house (for whatever reason). As far as compliance with (Grandma-Wide) goes, these routes are just as good as intending to visit your grandmother and intending to take the means you believe to be necessary (namely, driving to her house). And this will of course be true in general—for every end you don't intend and every means-end belief that you lack, as well as for every means that you intend, you will comply with a corresponding wide scope means-end requirement.

In contrast, like other narrow scope requirements, (Grandma-Narrow) only applies to you if you satisfy certain conditions—namely, if you intend to visit your grandma and believe that driving to her house is a necessary means of doing so. To not intend to visit your grandma ensures non-application, not compliance. Instead, the only way to actually *comply* with (Grandma-Narrow) is thus to intend to visit your grandmother and intend to take the means you believe to be necessary (namely, driving to her

house). Compliance involves *both* satisfying the conditions of application *and* doing what's required under those conditions.

This sort of dialectic can be rehearsed for each candidate requirement and its corresponding wide and narrow scope interpretations. Insofar as a given narrow scope interpretation delivers more intuitively plausible verdicts concerning compliance than the wide scope interpretation, this counts in its favor. Of course, such an argument is hardly decisive. Nonetheless, insofar as the notion of a requirement is supposed to be law-like—and nearly all major parties to the dispute agree that it is—then our judgments about compliance, non-application, and jurisdiction should carry dialectical weight.

Notice, though, that to address this concern the wide-scoper can help themselves to the distinction between jurisdiction and conditions of application and offer modified versions of the various wide scope requirements that do a better job jiving with judgments concerning compliance. For example, instead of (Grandma-Wide) they might offer something like (Grandma-Wide*):

(**Grandma-Wide***) (You intend to visit your grandmother ∧ believe that driving to her house is a necessary means to visiting her) —> rationality requires of you that ((you intend to visit your grandmother ∧ believe that driving to her house is a necessary means to visiting her) —> you intend to drive to her house).

Like (Grandma-Wide), we can interpret (Grandma-Wide*) as having jurisdiction over all rational agents, but unlike (Grandma-Wide) only applying to you if you satisfy certain conditions—namely, if you intend to visit your grandmother and believe you need to drive to her house in order to do so. If you do satisfy these conditions, then according to this interpretation of (Grandma-Wide*), rationality requires that you

intend to drive to your grandmother's house (in which case you'll be in compliance with the requirement) or else drop one of your antecedent attitudes (in which case the requirement will cease to apply). So it's not clear that the narrow-scoper has any real advantage over the wide-scoper when it comes to intuitions concerning compliance.[81]

However, there's another (relatively modest) advantage that the revised version of the narrow scope view enjoys over its predecessors. And that's that it helps blunt the force of one of the main objections to the narrow scope view—namely, that they seem false.[82] Consider (DC-Narrow), for example. It says that you are rationally required to believe *q* whenever you believe *p* and that *p* entails *q*. This might seem too strong. For aren't you at least sometimes *permitted* to give up one the antecedent attitudes, rather than believe *q*? (DC-Narrow) seems to assign a kind of fixed authority to the antecedent beliefs that is problematic; the agent should at least sometimes have the *option* of giving them up instead. And the same seems true of many (if not all) other narrow scope requirements.

The narrow-scoper I have in mind can meet the objector halfway. To begin with, we need to distinguish between being permitted to φ in the sense of *not being prohibited from* φ-ing (call this "weak permission") and being permitted to φ in the stronger sense of φ-ing being officially sanctioned ("strong permission").[83] The distinction is an important and intuitive one. For example, suppose a child finds some candy lying around the house. So long as no one in authority over her (such as her caregivers) has forbidden her from eating candy, she is permitted in the weak sense to eat it. If she does, she won't have

---

[81] Contra Lord (2014).

[82] [Citations: Broome, Gensler (?), Wallace, and others.] It's important to note, however, that the usual dialectic is complicated by the failure to adequately distinguish structural rationality from reasons rationality, and hence typically involves insufficiently discriminating claims about what one "rationally" ought to do or believe. The focus here solely concerns structural rationality.

[83] This is related but not identical to G. H. von Wright's (1970) distinction between 'weak' and 'strong' permission. Such distinctions figure prominently in both deontic logic and legal theory.

done anything she isn't supposed to do. But that might simply be because those in authority over her haven't (yet) considered the question of whether or not she should be allowed to eat candy. Now suppose they do consider the question and decide it's OK. This changes the normative landscape, and the child is now permitted in the stronger sense to eat candy. If she does she'll be doing something that enjoys the positive normative status of being sanctioned rather than just lacking the negative normative status of not being prohibited.

With this distinction in hand, the narrow-scoper can grant that narrow scope requirements like (DC-Narrow), even when they apply, "permit" you in the weak sense to give up one or more of the antecedent attitudes—they don't forbid you from revising (or otherwise ceasing to have) any of your existing beliefs. And so you might satisfy the antecedent conditions of (DC-Narrow)—or any other narrow scope requirement—at a particular time, and hence have it apply to you, and yet not do anything wrong by revising one of the relevant attitudes. In doing so you'll no longer satisfy the conditions of application, and the requirement will cease to apply. And if it doesn't apply, then you can't violate it. So in revising one or more of your antecedent attitudes you won't have done anything wrong by the lights of (DC-Narrow), since by doing so you'll have rendered it non-applicable, and hence silent.[84]

However, the narrow-scoper *shouldn't* say that narrow scope requirements like (DC-Narrow) "permit" you to give up one or more of the antecedent attitudes in the *stronger* sense of actually saying it's rationally OK to do so. One way of trying to make this permission explicit is the following:

---

[84] This is akin to Lord's (2011) notion of "exiting" a requirement, though he doesn't distinguish between weak and strong permission. Although the narrow scope view I'm considering is similar to Lord's, he (like Schroeder) fails to draw the distinction between conditions of application and jurisdiction. In more recent work (2014b), however, Lord does consider a similar distinction in response to an objection. However, the principles he's concerned with are diachronic rather than synchronic, which forces him to introduce additional complexities (such as an appeal to "cancelling conditions") that I wish to avoid. I'll briefly address the synchronic/diachronic issue below.

(**DC-Narrow***): (You believe $p$ ∧ you believe that $p$ entails $q$) —> rationality requires that (you believe $q$ ∨ not believe $p$ ∨ not believe that $p$ entails $q$).

But this would effectively turn (DC-Narrow*) into a wide scope requirement—it would be equivalent to (DC-Wide), except that it would only apply under certain conditions, just like (Grandma-Wide*) above. It would be equivalent to (DC-Wide*):

(**DC-Wide***): (You believe $p$ ∧ you believe that $p$ entails $q$) —> rationality requires that ((You believe $p$ ∧ you believe that $p$ entails $q$) —> you believe $q$).

(DC-Wide*) may very well be preferable to (DC-Wide), since it does a better job respecting our intuitions about compliance, understood per above, but it's not a narrow scope requirement—it requires that an agent (not) have a certain combination of attitudes, under certain conditions, and that's it.

Another way to make a "strong" permission to revise one's antecedent attitudes explicit in (DC-Narrow) would be as follows:

(**DC-Narrow****): (You believe $p$ ∧ you believe that $p$ entails $q$) —> rationality requires that you believe $q$ ∧ rationality permits you to (not believe $p$ ∨ not believe that $p$ entails $q$)

But (DC-Narrow**) is bizarre. It requires you to believe the consequences of other things you believe while simultaneously permitting you to revise those other beliefs. Why the differential treatment? It would be more plausible to change 'requires' to 'permits', but then the requirement would cease to be a

requirement. It would lose it's stringency, and so fail to do what it was supposed to do—namely, to explain why it's irrational to believe *p*, believe that *p* entails *q*, and fail to believe *q*.

## 2.5 The extent of requirements

The fourth question a theory of requirements faces is that of *extent*. That is, for each putative requirement *R*, there's the question: is *R local*, concerning a particular pattern of attitudes, or instead *global*, concerning the totality of one's mental state? Or is it more *regional*, concerning many but not all of one's attitudes? Requirements of rationality relate different possible attitudes that an agent might have (or lack), and the question of extent concerns *how many* possible attitudes a given requirement relates. The near-consensus is that many, if not all, requirements are local. Rather than taking into account all, or even most, of your mental states, they concern the relationships between a relatively small number of possible attitudes. The requirements we've considered so far are all local in this sense, and the default assumption is that all (or at least the vast majority) of requirements are similar in extent. This issue will be important in Chapter 5.

## 2.6 The temporal nature of requirements

The fifth question a theory of requirements must answer answer concerns their *temporal nature*. That is, for each putative requirement *R*, there's the question: is *R synchronic*, concerning what's required at a given time, or instead *diachronic*, concerning what's required over time? Historically, wide-scopers have tended to favor synchronic requirements whereas narrow-scopers have tended to favor diachronic requirements. The explanation for the former trend is rather straightforward, since it's hard to formulate diachronic wide scope requirements that are both (a) plausible and (b) not equivalent to their

synchronic counterparts. [FILL IN with example?] The explanation for the latter trend isn't as straightforward, and varies between narrow-scopers.

Kolodny, for example, thinks (at least some) coherence requirements should be analyzed as diachronic "process" requirements rather than synchronic "state" requirements. The primary difference between state and process requirement, according to Kolodny, is one of function: while state requirements "simply ban states in which one has conflicting attitudes", process requirements "say how, going forward, one is to form, retain, or revise one's attitudes" (2005: 517). That is, unlike process requirements, state requirements do not require you to "get out of" or to "get in" any particular state—they simply say something about "where you are" (2007a: 378).

There at two main reasons why Kolodny thinks process requirements are important. The first is that "ordinary attributions of irrationality are at least sometimes about what people do, or refuse to do, over time", and hence seem to concern rational transitions, or processes (2007a: 378). The second is that he aims to account for the thought that at least some requirements of rationality "can function as advice or guide your deliberation" (2007a: 378). This, he thinks is something only process requirements can do, since they tell you "to do something", whereas state requirements do not.

The first point is a fair one, but it's one that can be accommodated on a natural extension of the wide scope view—and it can be accommodated even more naturally on the pressure-based view, as will become clear in Chapter 5. Structural rationality, recall, is a matter of having the right 'structural' relations hold being one's attitudes. It's plausible, however, to think that structural rationality is not merely a matter of *having* the right kind of attitudes with the right kind of contents, but also a matter of the attitudes being *related* to each other in the right sort of way—that is, they need to be properly *based* on, or *responsive* to, each other. To intend to do something because you believe you ought to do it, for

example, seems appropriate in a way that not believing you ought to do something because you don't intend to do it does not—the latter is akin to wishful thinking or rationalization whereas the former just seems like a perfectly good transition. But the wide scope anti-akrasia requirement (AA-Wide) above does nothing to explain this apparent asymmetry—either transition is an equally good way of ensuring compliance, since (AA-Wide) only requires that you not *both* believe you ought to φ *and* not intend to φ.[85] Similarly, it seems perfectly appropriate to intend means *M* because you believe *M* is a necessary means to your end *E*, but to not believing *M* is a necessary means to your end *E* because you intend *E* and yet don't intend *M* seems terrible. Yet (ME-Wide) itself doesn't favor any particular route to compliance.

Broome takes our judgments concerning cases like these seriously, and so 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 (he calls these "basing permissions"). For example, to capture the apparent asymmetry between intending to do something because you believe you ought to do and not believing you ought to do something because you don't intend to do, Broome (2013: 141) offers the following basing prohibition (the label is mine):

(**AA-Basing**) Rationality requires of you that you not: not believe you ought to φ on the basis of your not intending to φ.

Broome himself offers (AA-Basing) as a diachronic principle—as governing the relationship between the two attitudes over time. This is because he takes basing to be a (at least partly) causal matter, and causal

[85] This is a version of the original "symmetry objection" advanced by Kolodny (2005), Schroeder (2009), and others. It's related but distinct from Lord's (2014) "real symmetry objection" concerned above.

relations are extended in time. This may not be a huge cost, but it does represent a departure from the core wide scope view, and it's not clear to me that it's necessary. A lot depends on how we understand the basing relation, of course, and like Broome (and everyone else) I don't have a worked-out theory to offer. Nonetheless, even if the process of forming attitudes on the basis of other attitudes takes time—as seems likely—once an attitude *A* is formed we still can ask, at any given time *t*, what *A* is based on at *t*. The answer will likely be complex, and indeterminate to some degree, but it's nonetheless plausible to expect one to be had. The fact that *A* is based on (say) *B* and *C* at *t* may be a partly historical (because casual) matter, but it's still a fact about *A* and about *t*.

Suppose, then, that we interpret (AA-Basing) synchronically, rather than diachronically. And suppose you find yourself in an akratic state: you believe you ought to φ but you lack the intention to φ. In order to be structurally rational, what do you have to do? Well, in order to comply with both (AA-Wide) and (AA-Basing)—which, as wide scope synchronic requirements, apply to you at all times—you must *either* form the intention to φ *or* drop the belief that you ought to φ on some basis *other than* your not intending to φ. What other basis for revision might there be? That depends on what other attitudes you have. And which are permissible? That depends on what other basing permissions and prohibitions there are. If you happen to believe that you lack sufficient reason to believe that you ought to φ, for example, then that might provide a suitable basis for revision of your belief that you ought to φ. But there's no guarantee that such an attitude will be present. And if you happen to lack any other attitudes that might serve as appropriate (i.e. permissible) bases for revision, then you'll be stuck—the only rationally available option will be to intend to φ.

The upshot is that although (AA-Wide) doesn't itself require you to "get in" a particular state, it may turn out that you're left with only one choice once your attitudes together with the various

(synchronic) basing principles are taken into account. The various structural requirements together with the basing principles may "corner" you, leaving you with only one rationally permissible option—only one "route" to compliance. In this way, a group of synchronic principles can have implications for what a person is rationally required to do over time—they can have dynamic upshot—without being formulated as diachronic or process requirements, as Kolodny would have it.[86] What "ordinary attributions of irrationality" about "what people do, or refuse to do, over time" show us is that structural rationality, and hence coherence, consists of more than just *having* the right sorts of attitudes—they also have to be had *in the right sort of way,* or held on an appropriate basis. (I'll return to this point in Chapter 5.) The upshot is that, like Kolodny, we should reject (State-Only):

> (**State-Only**) Structural rationality is exclusively concerned with certain patterns of) mental states, and not with their grounds or what they are based on.

But we don't yet have reason to reject (Synchronic):

> (**Synchronic**) Structural rationality is exclusively concerned with synchronic requirements.

I'm less sympathetic to Kolodny's second reason for focusing on process requirements: that only they can function as advice or guide your deliberation by telling "you to do something". First of all, it's not at all clear why we should think it's the job of the requirements of rationality to function as "advice" or a "guide for deliberation", nor is it clear what such a demand amounts to—both are dark phrases that Kolodny does little to illuminate. Indeed, although this particular criterion of adequacy is a recurring

[86] For recent defenses of synchronic (or "time-slice") theories of rationality in a broadly Bayesian framework, see Moss (2014), and Hedden (2015, forthcoming).

theme in Kolodny's work, it's not something he ever really argues for. Instead it's just assumed. This is puzzling because Broome—one of Kolodny's main targets—certainly doesn't assume they play such roles, and is quite clear they don't.[87] Secondly, it's worth noting that state requirements can, and do, tell "you to do something". So it's just not true that state requirements, whether construed as wide or narrow, fail to be response guiding. The *least* response guiding might be wide scope state requirements, since they do not require subjects to respond to a given conflict in one particular, privileged way, and instead leave it open the precise way in which the subject avoids being in an incoherent state. Nonetheless, wide scope requirements *do* require subjects to respond in some way. The requirement itself need not specify exactly what to do in order to be helpful. Consider practical advice. One might readily admit that specific, positive advice—e.g. "You should order lentil soup from Zaytoons for lunch today"—is often more helpful in guiding our deliberations than general, negative advice—"Don't consume harmful stuff"—while nonetheless insisting that they both qualify as genuine advice. The difference is less one of kind than of degree. In sum, Kolodny's case for process requirements is less than compelling; for all he says, (Synchronic) may well be true.

## 2.7 The normative status of requirements

The sixth and (for my purposes) final question a theory of requirements must answer concerns their *normative status*. That is, for each putative coherence requirement *R*, there's the question: in what sense, if any, ought we comply with *R*? This is a tricky question, of course, since 'ought' can be used in a number of different ways. (This is a point we'll return to in Chapter 3.) It's uncontroversial that coherence requirements—if such there be—are at least *weakly* normative insofar as they represent

[87] See Broome (2002, 2013) for his account of practical and theoretical reasoning and the relationship between reasoning and requirements.

standards of appraisal or evaluation, and for any such standard we'll be able to make claims about what you 'ought' to do relative to that standard. This is true of the law and etiquette just as much as prudence and morality. But coherence requirements are also plausibly thought to be normative in a stronger, less trivial sense. And that's because the charge of incoherence seems to involve a distinctive and genuine kind of criticism—one that targets you as an agent—that many other standards of appraisal do not. One needn't be subject to any kind of criticism merely for not doing what the law requires (after all, laws can be petty or unjust) or what some arbitrary standard of etiquette requires.

The recent debate over the normativity of structural rationality, however, is mostly concerned with the more specific question of whether we always have (at least some) *reason* to comply with structural requirements, or (more generally) why we should *care* about complying with them.[88] Not all sources of requirements are such that we always have reason to comply with them, after all. We do not always (if ever) have reason to do what Catholicism requires, for example, nor what the U.S. legal system requires. And although it's generally agreed that *being disposed* to be coherent is instrumentally valuable over the long run (at least for most people), nothing automatically seems to follow about the value of any particular manifestation of that disposition. Nonetheless, Broome (2013) is sympathetic with the thought that we always have reason to comply with the requirements of structural rationality, and actually endorses something stronger—namely, that the fact that structural rationality requires you to do something is *itself* a reason to do that thing. But he confesses that he has no argument for it. In contrast, Kolodny (2007: 241) finds it "outlandish that the kind of psychic tidiness that [any] requirement of formal coherence enjoins should be set alongside such final ends as pleasure, friendship, and knowledge". Of course, not all sources of reasons are on a par, so someone like Broome may well

[88] See, for exmaple, Kolodny (2005, 2008), Broome (2007, 2013), Southwood (2008), Way (2009), and Raz (2011), among many others.

agree with the sentiment expressed by Kolodny and nonetheless think of coherence requirements as sources of less-than-weighty reasons.

Suffice it to say, there's a large (and sometimes murky) debate concerning the normativity of structural rationality, and it's one I won't enter into in any depth here. But it's worth mentioning because it's a question—like those above—that any worked-out theory of requirements has to address.

Although each of the issues I've canvassed in this chapter will be relevant at different points in what follows, it is the question of "scope" that has received the most attention, and it's the one that I'll be focusing on in Chapters 3 and 4. I'll mentioned the others when relevant, and return to several of them in the Chapter 5 when comparing the relative merits of the requirements-based picture and the pressure-based one.

# CHAPTER 3: CONDITIONALS AND REQUIREMENTS

In this chapter I present and criticize a historically influential argument in favor of the wide scope view. It's a broadly linguistic argument, involving an appeal to our judgments involving ordinary deontic conditionals (or "iffy oughts"), and it's a bad one. It turns out the linguistic data is basically *silent* on the relative merits of the wide scope and narrow scope views. Purely philosophical considerations will have to carry the day, and on that score I don't think there are any decisive considerations favoring either the wide or the narrow scope view.

## 3.1 Conditionals and conditional requirements

At the center of the debate over the nature and existence of structural requirements has been the question of how we should understand their *scope*—that is, what *form* they take. As it turns out however, there are (at least) two importantly different issues that have traditionally been conflated. The first concerns whether structural rationality ever requires one to have particular attitudes, or instead is exclusively concerned with mandating or prohibiting combinations of attitudes. The second concerns the proper interpretation of certain natural language conditionals containing modal expressions like 'ought' and 'requires'—conditionals that have traditionally been taken to be capable of expressing (potential) requirements of rationality. It's the former issue that I take to be of more fundamental importance, though it's the latter issue that has received the most attention in debates over the nature of structural requirements—mainly because it has (wrongly) been assumed to be relevant to the former.

It's an easy mistake to make. And that's because it's very natural to (try to) express structural requirements using conditionals. Structural rationality, after all, is fundamentally a matter of how one's attitudes relate to each other—of how they "fit" or "hang" together—and conditionals give us a natural way of expressing claims about such relations. In particular, conditionals allow us to express claims about which attitudes are rationally required *given certain other attitudes*, and thereby give voice to our intuitive judgments of proper and improper fit between them—the very judgments that prompt theorizing about structural rationality in the first place. What's more, it appears to be a truism that if rationality *requires* you to φ then you rationally *ought* to φ, so not much seems to hinge on whether we use 'requires' or instead 'ought' in formulating our requirements.

Given all this, it's understandable that much of the debate over the nature and scope of structural requirements has focused on the proper interpretation of claims like the following:[89]

| | |
|---|---|
| **Deductive Coherence** | If you believe *p* and that *p* entails *q*, then you rationally ought (or are required) to believe *q* |
| **Strict Means-End** | If you intend *E* and believe that *M* is a necessary means to *E*, then you rationally ought (or are required) to intend *M*. |
| **Anti-Akrasia** | If you believe you ought to φ, then you rationally ought (or are required) to intend to φ. |

[89] Minor modifications can be (and oftentimes are) made to increase plausibility, such as a clause in Deductive Coherence stipulating that *q* isn't completely irrelevant (cf. Harman (1986)) or a clause in Strict Means-End stipulating that you don't believe you will ψ regardless of whether you intend to ψ. These sorts of qualifications won't be of direct concern in what follows.

Minor quibbles about content aside, these claims enjoy considerable intuitive support. And at least at first blush they seem to hold in full generality—we can add 'necessarily' to the front of each without sacrificing much, if any, plausibility. As a result, such conditionals have commonly been taken to express —or at least be capable of expressing—requirements of rationality, understood as strict, exception-less principles applying all rational creatures (or at least those satisfying the antecedent).[90] It's (schematic) requirements such as these that are then thought to underwrite our judgments of (ir)rationatity in particular cases, as well as particular types of cases. Recall Tom, for example. He believes that he's Superman and that being Superman entails being able to fly, but fails to believe that he (himself) can fly. He thus violates Deductive Coherence: he satisfies the antecedent but doesn't have the further attitude that he rationally ought to have—namely, the belief that he can fly.

Despite the evident plausibility of claims like those above, however, it's standardly thought there's a problem: the requirements expressed seem subject to counterexample.[91] You might believe *p* and that *p* entails *q,* for instance, without it thereby being the case that you rationally ought to believe *q*. After all, *q* might be wildly implausible, and you might recognize it as such, in which case it's not true that you rationally ought to believe *q*. Indeed, it seems strange to think that structural rationality would even *permit* you to believe something you took to be (literally) incredible. In such a case although it seems clear that you *do* need to revise one of your beliefs in order to avoid incoherence, it doesn't seem

---

[90] Strictly speaking what would be expressed are requirement *schemas*; particular requirements would only be expressed once the appropriate substitutions are made.

[91] See Broome (1999, 2013), among others, for this worry. It's important to note, however, that the usual dialectic is complicated by the failure to adequately distinguish structural rationality from reasons rationality, and hence typically involves insufficiently discriminating claims about what one "rationally" ought to do or believe. The examples here focus on conflicts between one's attitudes, and hence concern structural rationality. They therefore differ from the examples from Chapter 1, which focused on apparent "conflicts" between structural rationality and reasons rationality.

there's any particular attitude that's singled out for revision—contra what Deductive Coherence says. Strict Means-End seems subject to similar counterexamples. For suppose you intend *E* and believe that *M* is a necessary means to *E*, but also believe that taking the means *M* would have terrible consequences for your friends and family. Is it really true that you rationally ought to intend *M* anyways? Arguably not. Like Deductive Coherence, Strict Means-End appears to attribute a kind of authority to your antecedent attitudes they don't in fact have—at least not in general. As might be expected, this problem generalizes: for many if not all claims like Deductive Coherence and Strict Means-End, there will be cases in which one has the attitudes specified in the antecedent and yet the normative claim specified in the consequent —namely, that one *ought*, or is *required*, to have some further attitude—is too strong. Call this the "too strong problem".

There's another problem in the vicinity. For suppose that you happen to be in the normatively conflicted state of believing that you ought to φ and that you ought not to φ. Does it really follow that you rationally ought to intend to φ *and* intend not to φ, and thus have inconsistent intentions? Arguably not—it's doubtful that structural rationality would conflict with itself in such a way, contra what we might expect if Anti-Akrasia were true.[92] Or suppose you believe *p* and that *p* entails *q*, but also believe *r* and that *r* entails not-*q*. Given Deductive Coherence, it seems to follow that you rationally ought to both believe *q and* believe not-*q*—and thus have inconsistent beliefs. But, again, this seems implausible. Structural rationality is supposed to *prohibit* incoherence, not require it. Call this the "conflict problem".

In response to these apparent problems, there's a standard solution. As Broome (2007) notes:

> English sentences [do] not always give modal operators their logically correct scope. Compare the sentence 'If you jump from a plane, you must have a parachute to survive'. This really means

[92] Cf. Broome (2007, 2013).

> that you must (if you jump from a plane, have a parachute to survive). The scope of the 'must' really includes the whole conditional, even though in the literal English it includes only the consequent. (10)

Thus, deontic conditionals like Deductive Coherence, Strict Means-End, and Anti-Akrasia (as well as their instances) are standardly taken to be *scope ambiguous*, having one interpretation which (at least at some level of analysis) can be formally represented as

**Wide** $O(C1 \Rightarrow C2)$

and the other of which can be formally represented as

**Narrow** $C1 \Rightarrow O(C2)$

where '⇒' is a two-place conditional operator, 'O' is a one-place modal operator (representing the contribution of 'ought', 'requires', etc.), and 'C1' and 'C2' are the relevant conditions (e.g. an agent's having certain attitudes or reason(s)).[93,94]

---

[93] For versions of this claim, see Broome (2013), Brunero (2010), Dancy (1977), Greenspan (1975), Gensler (1985), and many others—including various Medieval philosophers. Schroeder (2004, 2011b) notes the widespread tendency to posit ambiguity, but he resists the trend by arguing that the relevant 'ought' is not a sentential operator and so is incapable of entering into the relevant scope relations. As we'll see, Schroeder is right to deny the ambiguity, but he's right for the wrong reasons—the ambiguity claim rests on an implausible view of 'if', not 'ought'. For a development of the standard "flexible contextualist" account of 'ought' and other modals that can accommodate Schroeder's data, see Hacquard (2010) and Kratzer (2012).

[94] As the quote from Broome indicates, the fact that both the Wide and Narrow interpretations of conditionals like Deductive Coherence are available—and that the Wide interpretation is to be preferred—appears to be independently motivated. And that's because there are many other conditionals containing modals that similarly seem to resist "detaching". [Citations and examples?] In each case the same strategy is available insofar as we want to vindicate the intuitive acceptability of the original conditional while avoiding the counterintuitive consequences: the modal operator can be interpreted as taking wide scope relative to the conditional itself. The fact that this maneuver appears to be available elsewhere provides support for its use here. But it turns out such a maneuver is implausible across the board.

In this way, many philosophers argue that the surface structure of ordinary language is misleading, and that the deontic conditionals like Deductive Coherence, Strict Means-End, and Anti-Akrasia should be given a non-obvious logical form, with the modal operator taking "wide scope" relative to the conditional operator rather than "narrow scope" (hence the names). And that's because only the Wide interpretation avoids counterexamples like those above. It does so in virtue of resisting so-called "factual detachment": given C1, it follows from Narrow but not from Wide that O(C2). On the Wide interpretation, then, all that that you ought (or are required) to do is make the conditional true (or at least not false)—you do *not* have to have the particular attitude specified in the consequent.

Notice that this broadly linguistic argument in favor of the Wide interpretation bears quite directly on question of "scope" from Chapter 2 concerning whether structural rationality ever requires one to (not) have particular attitudes, rather than merely *combinations* of attitudes. For if the structural requirements expressed by the conditionals above should be interpreted along the lines of Wide rather than Narrow, then they'll only be concerned with combinations of attitudes rather than individual ones. And since it's plausible that the same dialectic arising with Deductive Coherence, Strict Means-End, and Anti-Akrasia will arise with respect to (nearly) all other candidate requirements, the seemingly inevitable conclusion is that we should be wide-scopers across the board—all structural requirements will be wide scope. The broadly linguistic considerations adduced above have thus played an important role historically in motivating the wide scope view, even though the wide scope view *itself* shouldn't be understood as a linguistic claim.

Before proceeding to evaluate the wide scope view and the linguistic argument in favor of it, it's worth noting one additional benefit it enjoys. And that's that it opens up the possibility that the difference between reasons rationality and structural rationality can be explained, at least in part, by

differences in the formal structure of the requirements corresponding to each. For narrow scope interpretations of sentences expressing (at least apparent) "requirements" of reasons rationality seem perfectly acceptable:

**Conclusive Reason** If you have conclusive reason to believe $p$ and that $p$ entails $q$, then you rationally ought (or are required) to believe $q$.

**Decisive Reason** If you have decisive reason to $\varphi$, then you rationally ought (or are required) to (intend to) $\varphi$.

Unlike structural requirements, then, (purported) requirements of reasons rationality unproblematically require something specific of you—they single out a particular attitude as favored. They are not neutral, or symmetric in terms of compliance—there's only one way to satisfy them.

The same general point can be made even if, as seems likely, it's a mistake to think of reasons rationality as fundamentally requirements-based in the same way that structural rationality is. For unlike structural rationality, it's common to think of the normative notion(s) at the heart of reasons rationality —namely, reason(s), justification, evidence, etc.—as coming in degrees, or as being "weighted", and as being normatively non-strict. Yet neither of these is characteristic of a requirement, which is a strict, all-or-nothing affair—a point I'll return in Chapter 4. And so even if reasons rationality is not requirements-based, so long as structural rationality *is* requirement-based, then we'll still be draw the distinction between dimensions (at least partly) in terms of differences in the nature or form of the normative phenomena characteristic of each. The popular wide scope view of structural rationality might therefore seem to offer a promising account not only of how we should understand the nature of structural

rationality, but also of what distinguishes it from reasons rationality. It's no surprise, then, that the view has been adopted by so many philosophers in recent decades.[95]

Despite its initial plausibility, we should be reluctant to embrace the wide scope view. But that doesn't mean we should opt for the narrow scope view instead. That's because we can—and, I'll argue, should—reject the terms of the debate, and thereby reject both. There are two assumptions undergirding the recent scope debate that are especially problematic. The first is that conditionals like Deductive Coherence are ambiguous between Wide and Narrow interpretations. The second is that structural rationality is fundamentally requirements-based. Both assumptions are false. The first problem isn't fatal, and leaves the core idea(s) of the wide scope view intact, but it's worth noting because it undermines the historically prominent and influential line of argument in favor of the wide scope view outlined above. It's also independently important, since all parties to the debate need to be able to make sense of our ordinary thought and talk about rationality. Accordingly, one of the main aims of this chapter is to get clear on what readings of deontic conditionals like Deductive Coherence are actually available. The second problem, however, cuts to the core. It'll be the focus of Chapter 4.

## 3.2 Modals

The first problem with the wide scope view is one of motivation: the linguistic—and historically very influential—argument in support of it rehearsed above isn't a good one. That doesn't show that the wide scope view is false, of course. But it does undermine one of it's main pillars of support, at least historically.

---

[95] See, for example, Broome (1999, 2005), Brunero (2010), Darwall (1983), Greenspan (1970), Hill (1973), Ross (2012), Shpall (2013), Wallace (2001), Way (2010, 2013), Wedwood (2011).

The linguistic argument in favor of the wide-scope view rests on three assumptions. The first is that conditionals like Deductive Coherence and Strict Means-End have at least one reading on which they are true. The second is that they are ambiguous. The third is that they are ambiguous between Wide and Narrow readings. While I agree with the first two assumptions, I reject the third: although such conditionals do have at least one reading on which they are true, and are in fact multiply ambiguous, they are not ambiguous between Wide and Narrow readings. All three assumptions are important. The first is important because debates over structural requirements are awash in iffy oughts like Deductive Coherence and Strict Means-End, and it's intuitive judgments concerning them that provide the impetus for much of our theorizing. So it's worth getting clear on which reading(s) are available, and this is directly related to the second assumption. But it's the third (and faulty) assumption that matters most for present purposes. I'll return to the second—and explain why it's correct—later in this chapter, and I'll return to the first in Chapter 5. Once the positive view is on the table I'll be able to offer an an alternative, non-requirements-based explanation of why conditional sentences like Deductive Coherence and Strict Means-End have at least one reading on which they are true.

As we've seen, the debate over the logical form of propositions expressed by deontic conditionals (or "iffy oughts") like Deductive Coherence or Strict Means-End concerns the relationship between two things: the 'if' and the 'ought'. It's time to say more about each. Although the main objection to the Scope debate concerns its treatment of 'if', for reasons that will soon become clear, I first need to say something about 'ought' and related modal expressions.[96]

So far I've been not-so-tacitly assuming that modals like 'ought', 'should', and 'requires' can be used in more than one way. (For concreteness, and in deference to tradition, I'll continue focusing on

[96] The discussion of modals and conditionals below is especially indebted to Geurts (2004), Hacquard (2010), Knobe and Szabo (2013), Khoo (2013), Kratzer (2012), and Rothschild (forthcoming).

'ought'.) This isn't particularly controversial. Indeed it's common among philosophers to distinguish different "senses" or "uses" of 'ought', where this is standardly taken to include (at least) the *moral* 'ought', the *legal* 'ought', the *prudential* 'ought', and the *rational* 'ought'. It's also standard to distinguish between "pro tanto" and "all things considered" uses of 'ought'. (This point which will be important later on.) Philosophers often assume that the possibility (and necessity) of drawing such distinctions shows that 'ought' and the like are ambiguous. If such a claim is merely meant to highlight the fact that 'ought' admits of multiple, non-equivalent readings, it's harmless enough; but if it's meant as a semantic claim, it's false. Modal expressions like 'ought' are *not* semantically ambiguous—at least not according to linguistic orthodoxy. Instead, they are context-sensitive.

To see why that is, I'll introduce and motivate a simplified (and highly informal) version of the standard "flexible contextualist" account of modals introduced by Angelika Kratzer (1977, 1981), and developed by many others since. Doing so will allow us to better understand what's going on with deontic conditionals like Deductive Coherence and Strict Means-End, as well as a range of other claims made throughout this dissertation. The main reason I'm focusing on Kratzer's account of modals (and conditionals) is because I'm convinced that it's the *correct* account, at least in broad outline.[97] Although various challenges have been raised in recent years—many by philosophers—none of them, to my mind, reveal the need for radical revisions to the core ideas; at most they reveal the need for minor modifications, as well as the need for supplementation by (and integration with) theories of other

[97] One particularly important challenge to the standard Kratzerian account arises from its treatment of "graded" modality and modal expressions. Lassiter (2011) develops the challenge at length and offers a probability-based alternative. (See also Yalcin (2010), among others.) For a defense of Kratzer-friendly account of graded modality, as well as criticism of Lassiter, see Klecha (2012, 2014). Indeed, as Klecha points out, just as there are both gradable and non-gradable adjectives, as well as gradable and non-gradable nouns, so we shouldn't be surprised to find gradable and non-gradable modals. So it's a mistake to assume that all modal expressions belong in one category or the other. For an accessible overview of various issues arising with graded modality, see Lassiter (forthcoming) as well as Portner and Rubinstein (forthcoming).

semantically and pragmatically important phenomena. A fully satisfactory account of modals, for example, needs to explain the complex interactions between modals and tense and aspect, as well as between modals and a number of different phenomena that fall under (what we might call) "discourse pragmatics".[98] Though important, none of these complexities will be dealt with here. The simplified version will suffice for my purposes.

Although my focus has been (and will continue to be) on modal verbs like 'ought', it should be noted that modal expressions belong to a variety of different syntactic categories, all of which are standardly interoperated as taking clauses for complements:[99]

| | |
|---|---|
| Modal verbs: | *have (to), ought (to), need (to), require (to),* etc. |
| Modal auxiliaries: | *must, can, should, might,* etc. |
| Modal adjectives: | *possible, necessary, probable,* etc. |
| Modal adverbs: | *possibly, necessarily, probably, maybe,* etc. |

One of the most important features of such modal expressions, considered as a class, is that they are *flexible*, in the sense that they can be used in different ways in different contexts (and oftentimes within the same context). Consider, for example, the following sentences involving 'have (to)':

---

[98] The literature is vast and rapidly growing. For a book-length overview of the central issues, see Portner (2009). For some recent work defending a broadly Kratzerian account of modals in the face of recent challenges, as well as attempts to integrate it with other semantically and pragmatically important phenomena, see Hacquard (2010, 2013a, 2013b, forthcoming), Dowell (2011, 2012, 2013, forthcoming), Arregui (2010, 2011), Klecha (2014), and Silk (forthcoming, book ms).

[99] Clauses are usually taken to express propositions; this is why modals are generally modeled as propositional operators.

(1) You have to take the F train.

(2) I have to sneeze.

(3) You have to pay me back.

In the absence of special information, it would be most natural to expect (1) to be used to make a claim about what is necessary in order to (optimally) achieve certain aims or goals, and thus to express a claim involving so-called *teleological necessity*. In contrast, it would be most natural to expect (2) to be used to make a claim about what is necessary given your current physical state, and thus express a claim involving so-called *circumstantial necessity*. Whereas with (3) it would be most natural to expect it to be used to make a claim about (some kind of) obligation, and hence to express a claim involving so-called *deontic necessity*. Each use thus corresponds to a distinct kind of modal "flavor". There are a variety of means by which we can make the intended flavor of the modal more explicit. For instance, we might use a *given*-clause:

(1') Given your goals, you have to take the F train.

(2') Given my current state, I have to sneeze.

Or we might use an adverb:

(3') Morally, you have to pay me back.

Again, there are a variety of linguistic means of achieving the same end, though oftentimes they won't be required—it'll be clear enough in context what the intended interpretation is.

As the informal glosses of (1)-(3) above suggest, at the most fundamental level, what modal expressions do is relate the information provided by their syntactic complement, or "prejacent"—e.g. the proposition expressed by 'you take the F train' in (1)—to a distinct, contextually provided body of background information, where this information represents (some combination of) what is known, what the circumstances are, what is desired, what is aimed at, what morality commands, etc.[100] (Notice that information, in the sense at issue, needn't be *true* information—hopes, dreams, works of fiction, and political speeches contain lots of information, very little of which corresponds to anything in reality.) Modal expressions thus function as modal operators—they express relations between two (possibly complex) pieces of information, one of which is provided by the prejacent and the other of which is (for the most part) backgrounded and provided by context. The modal's flavor is determined by the latter—that is, by the kind of information contained in the background, and hence being "held fixed" for the purposes of evaluating the modal claim.

As already noted, the availability of multiple readings or interpretations of modals like 'have (to)' does not mean they are ambiguous. Indeed, the systematicity and cross-linguistic regularity of available readings counts rather decisively against an ambiguity theory. Instead, what it reveals is that modals are context-sensitive, with their "flavor" being fixed by some feature of the context.[101] The implausibility of an ambiguity theory is made even more manifest once it's recognized that each flavor of

---

[100] This is compatible with modal expressions doing (or being used to do) other things, of course, in addition to this.

[101] Here and throughout I'm using 'context' broadly so as to include the communicative intentions of the speaker. Though I myself have strong neo-Gricean sympathies, I don't want to take a stand on the question of what 'context' amounts to; instead I'll cheat just like everyone else.

modality divides further into various "sub-flavors". I've already made this point with respect to 'ought'—it obviously admits of a range of deontic sub-flavors, including those relating to morality, the law, prudence, and rationality. But 'ought' is hardly unique in this respect. Take (3), for example, on its deontic reading. In addition to being used to state what you're morally obligated to do, as originally envisioned, (3) can instead be used to make a claim about what it would be *in your best interest* to do (where presumably this isn't just a function of fulfilling your moral obligations), or what it would be best for you to do given *my* interests. Again, the different readings can be made explicit in various ways, though (depending on the context) they needn't be:

(3'') Given your interests, you have to pay me back.

(3''') Given what I want, you have to pay me back.

Or if we were to assume that you lost a court battle, one might use (3) simply to state what the law requires. It would be equivalent to saying something like:

(3'''') Legally, you have to pay me back.

Other standard (sub-)flavors include: epistemic (roughly, concerning what is known), doxastic (concerning what is believed), bouletic (concerning what is desired), and stereotypical (concerning what

is typical or normal).[102] And of course, if the opening reflections of this dissertation are not wide of the mark, we should recognize not just one but *two* rational "flavors" (or sub-flavors, or sub-sub-flavors, depending on one's preferred classification)—one corresponding to reasons rationality and the other to structural rationality.

Many philosophers will also insist on there being distinctive logical, conceptual, physical, alethic, and/or metaphysical modalities, and that these are expressible in natural language in one way or another. In general, we should expect there to be a distinct "flavor" corresponding to each sense of "possibility" and "necessity", as well as to each dimension of evaluation that we're sensitive to. Of course not all modals admit of each flavor—far from it. While many modals can express a variety of different kinds of flavors, others are more restricted in the kinds of interpretation they can receive. Some restrictions are due to grammatical constraints, but others appear to be lexically specified. Although 'have (to)' is rather promiscuous, for example, 'must' is less so and 'might' is much less so—'might' is almost exclusively epistemic.[103] The same pattern holds across languages.[104]

It's of course a hard question how we should understand the nature or source of each flavor, and how (if at all) they're related to each other. As Kaufmann and Kaufmann (2015) note, formal semanticists tend not to pay careful attention to taxonomic considerations, and usually start out with a

---

[102] The multiplicity of possible flavors is of course compatible with certain sentences have a "default", or otherwise favored, reading. This will oftentimes be a function of the kind of context the sentence is most likely to be used in, and hence the status of a certain reading as favored will typically be pragmatic and capable of being cancelled or overridden. For example, in suitable conversational contexts (1) above —'You have to take the F train'—can be used to express a deontic, rather than teleological, claim, and (2)—'I have to sneeze'—can be used to express a teleological, rather than circumstantial, one.

[103] A point often overlooked in the linguistics literature is that 'might' also has (what appears to be) a purely metaphysical interpretation. This is natural interpretation of 'might' in contexts where we're asked—typically by a philosopher—to consider various different (metaphysically possible) ways "the world" might have been.

[104] Cf. Traugott (1988), Bybee, Perkins and Pagliuca (1994), and Palmer (2001), among others. Nauze (2008), however, provides evidence that this cross-linguistic pattern isn't as robust as originally thought.

"coarsely grained labelling which is only refined further if particular characteristics turn out to be relevant for the analysis at hand". This is one area (among others) where I think philosophers are particularly well-suited to make contributions to the existing literature in linguistic semantics. The need to distinguish between two distinct "rational" flavors of 'ought' (etc.) can be seen as one (fairly modest) example of how further progress can be made.

But differences in flavor is only part of the story. Modals also differ in their *force*, or strength. Differences in modal force correspond to differences in the *kind* of relation expressed between the relevant body of background information and (the propositions expressed by) their prejacent. Unlike flavor, force is always lexically specified, at least in English. For example, some modals, such as 'have (to)' and 'must', are "necessity" modals, while others, such as 'possibly' and 'may', are "possibility" modals. Very roughly, necessity modals express relations of *consequence* or *entailment* between their prejacent and background information, while possibility modals express relations of *compatibility* or *consistency* between their prejacent and background information. To say, as in (2'), that that you *have to* sneeze, given your circumstances, is (very roughly) to say that your sneezing is a (physical) *consequence* of your current state, whereas to say that you *can* sneeze is to say that your sneezing is *compatible* with your current state. Similarly, to say, as in (3'), that you *have to* pay me back, morally speaking, is (very roughly) to say that your paying me back is *entailed by* the demands of morality, whereas to say that you *may* pay me back is to say that your paying me back is *consistent with* the demands of morality.

There is controversy over the exact nature and appropriate formal representation of modal flavors and forces, though the centrality of these two parameters of lexical variation is largely unquestioned. Traditionally, however, sentences containing modals have been given quantificational truth conditions, with modals themselves (at least in their verbal form) being treated as quantifiers over

possibilities. For the purposes of this dissertation, I'll be following suit. To be clear: this is a theoretical choice. I said above that modal expressions function as modal operators, expressing relations between two (possibly complex) pieces of information—one provided by the prejacent and one provided largely by context. This was meant as a definition, not an analysis. How do we get from there to an analysis of modals as *quantifiers*? By making two popular, but ultimately optional, theoretical choices. The first is to follow Kripke (1959) and company in analyzing pieces of information (or "propositions") as sets of possibilities. Possible worlds are of course one (rather maximal) type of possibility, but many philosophers and linguists take possibilities to be more fine-grained, and reliance on worlds is replaced by reliance on situations, events, cases, actions, and the like, or some combination thereof. For my purposes it won't be necessary to take a stand on what possibilities amount to exactly. I'll follow Kratzer (2012) in being maximally permissive. The second theoretical decision is to follow Barwise and Cooper (1981), among others, in analyzing quantifiers as relations between sets, in accordance with Generalized Quantifier Theory. We thus arrive first at an analysis of modal operators as relations between two sets of possibilities, and then at an analysis of modals as quantifiers over possibilities.

However, whether modal expressions should, in the final analysis, be analyzed as quantifiers at the level of logical form is not my concern. My main concern is instead with the truth conditions of modal claims, and it's at *that* level that I find the quantificational account compelling. So although I'll be assuming that sentences containing modals have quantificational truth conditions, and that modals themselves are a type of quantifier, I'm more attached to the former than the latter.

Differences in force and flavor are easily captured on the quantificational account of modals. Modal flavor remains a function of the modal's informational background, with the latter now being understood to furnish a set of (relevant) possibilities that serves as the modal's domain. As before, the

set of possibilities (and hence flavor) can vary from context to context. Differences in modal force now amount to differences in *quantificational* force, with possibility modals functioning as existential quantifiers, requiring that their prejacent be true in *some* of the (relevant) possibilities, and necessity modals functioning as universal quantifiers, requiring that their prejacent be true in *all* of them.[105] There are thus three elements to any modal construction: the quantifier, the prejacent, and the domain of relevant possibilities.

Which possibilities are relevant? Again, it depends. Just as with other, more familiar quantificational expressions, such as nominal quantifiers ('all', 'some', 'most'), the relevant domain will typically be restricted in some (possibly quite severe) way. Since quantification in natural language almost always involves quantification over a contextually restricted domain, it shouldn't be surprising that modals behave similarly, given that modal claims are quantificational claims.

One of Kratzer's key contributions concerns the mechanics of modal domains—of how they are constructed—and hence what determines a modal's flavor. In particular, she has argued that the modal domain arises from the interaction between the values of two distinct contextual parameters, which she calls the "modal base" and the "ordering source", both of which can be given a formally clear and well-understood representation.[106] However, the details of her particular implementation aren't important for our (rather informal) purposes. All that matters is that, according to the standard view, there are two things we need to know in order to evaluate a modal claim: first, which possibilities are relevant and,

---

[105] In what follows I'll mostly be ignoring differences between s-called "weak" necessity modals (e.g. 'ought' or 'should') and "strong" necessity modals (e.g. 'must' or 'have to'). (For a recent account of the difference, as well as further references to the literature, see Rubinstein (2014).) I'll also be ignoring non-modal aspects of meaning of modal expressions, such as their possible status as "evidentials" (cf. Willett (1998), Faller (2002), and Murray (2010), among others) as well as possible performative uses (cf. Ninan (2006)).

[106] Again, see Kratzer (2012) for details, as well as myriad of work by others. For relatively accessible overviews, see von Fintel (2006), Portner (2009), and Hacquard (2011).

second, which particular dimension of evaluation the possibilities are to be compared or ranked on the basis of. In short, we need (what I'll call) a *modal background* and a *ranking*. Together these determine the modal's domain, which is the set of "highest ranked" possibilities that are relevant. In this way, modals are taken to quantify over a more restricted set of possibilities (i.e. the highly ranked ones) than the modal background itself provides.[107]

One of the strengths of the classic Kratzer framework is it's flexibility.[108] Indeed, it's flexibility extends beyond what is commonly supposed—beyond, that is, the flexibility it allows for in determining which possibilities are relevant and which flavor (or ranking) is operative. In particular, the framework

---

[107] My informal gloss of the semantics of modals departs from Kratzer's in two (relatively minor) ways. First, while I'm working with (what's known as) an "ordering" semantics for modals, Kratzer prefers a "premise" semantics. On her view, both contextual parameters operate in the same way: they are (functions from possible worlds to) sets of propositions, or "premise sets". So rather than work directly with a set of possibilities and a ranking, she works with two sets and extracts a ranking from one of them. The difference is mainly a technical one, however, and as Lewis (1981) points out, a theory framed in terms of the one can be translated into the terms of the other without remainder. I prefer to operate with an ordering semantics because I find it more intuitive, but nothing hangs on the decision, technically speaking. Second, unlike Kratzer, I'm taking for granted the so-called "limit assumption", according to which there's always a non-empty set of most highly ranked (or "best") possibilities. To avoid unwanted predictions in cases where the limit assumption is violated (i.e. in which there is no set of "best" possibilities) Kratzer complicates the semantics by permitting infinite chains of relevant possibilities that are ranked progressively higher without limit by the ranking. The limit assumption remains controversial, and there are attractive alternatives that represent a halfway house between wholehearted embrace of the limit assumption and wholesale abandonment of it. Swanson (2011), for example, defends the limit assumption and argues that apparent exceptions can be handled by a non-semantic mechanism. But regardless, since violations of the limit assumption only arise in cases involving infinite domains, the simpler semantic picture I adopt will suffice.

[108] Somewhat ironically, this is also considered by some to be one of its most serious weaknesses. In many cases the range of admissible interpretations is far more constrained than the bare-bones Kratzerian account would predict. But, as noted above, all this shows is that the bare-bones account needs to be supplemented by, and integrated with, insights and theories from other areas of study, both within linguistics and beyond. There's a division of labor among different (sub-)branches of linguistics, as well as between those of linguistics and those of philosophy, psychology, cognitive science, and the like. This division of labor facilitates progress, but there's inevitably a certain amount of arbitrariness in how things are divided, as evidenced by the advent of sociolinguistics, psycholinguistics, cognitive linguistics, semiotics, formal pragmatics, metasemantics, and the like. A fully mature theory will be an integrated one.

allows for the possibility of what Knobe and Szabo (2013) call *impure modals*. They reject the assumption of modal purity, understood as the claim that in any given context a modal has at most one flavor (circumstantial, teleological, deontic, etc.) and never involves a "mix" of flavors. As they point out, while the assumption of modal purity might be attractive on a simple ambiguity theory—one that posited a list of possible flavors and individuated modals in terms of them—it would appear to be dubious on a flexible contextualist theory. As they put it:

> The most natural way of thinking about conversational contexts would be to assume that they embody a mixture of different information. In a given context, we might be primarily concerned with the circumstances but also somewhat concerned with not acting immorally and achieving certain goals. In other words, when one shifts…to a theory based on context, it is only natural to suppose that there can be *impure modals*. (25)

The possibility of certain kinds of modal impurity should hardly be surprising. Philosophers have long distinguished between different dimensions of (broadly) normative evaluation. We realize, for example, that what we *morally* ought to do can come apart from *prudence* or *self-interest* might tell us to do. In much the same way, what we *epistemically* ought to believe might very well come apart from what we *pragmatically* ought to believe. However, according to many (but not all) philosophers, settling each of these questions on their own doesn't settle the question of what we *all things considered* ought to do. That depends on how the different dimensions or sources of normativity "stack up" or get "ranked" in the final analysis. Perhaps morality always trumps when it's in the contest; perhaps not. The point is just that the"all things considered" question involves a form of impurity—*deontic* or *normative* impurity.

We're also obviously comfortable mixing normative and (broadly) probabilistic information in ordering or ranking alternatives, as we do whenever we're engaging in expected utility calculations.

What Knobe and Szabo are pointing is just that we should also expect *further* impurity. Of course, just because some modals allow for impurity it doesn't follow that all modals do, nor does it follow from the fact that impurity is possible that impurity is present in a given case. Indeed, there are several ways of encouraging and (at the limit) enforcing purity of flavor. One is simply in the choice of modal expression used, since some modals are subject to more severe lexical restrictions on the kind of flavor or sub-flavor they're allowed to have than others (e.g., 'might' vs. 'have (to)') and many modals come paired with a preferred , but not required, flavor (e.g., 'might' vs. 'ought'). But, as noted already, even if one uses a modal expression that admits of more than one (sub-)flavor, such as 'ought', one can make the intended (sub-)flavor explicit by, for example, a using a suitable adverbial modifier—one that makes clear the (at least dominant) basis of the ranking—such as 'morally', 'practically', 'rationally', or whatever.[109]

Lexical restrictions, on the present view, are to be understood as constraints on the type of modal background and ranking that a given modal expression is allowed to have. And it may be some some lexical restrictions require the ranking associated with a given modal to be pure. But this assumption is optional. For as Knobe and Szabo point out, all that's needed is the significantly weaker assumption that "the outer domain and the ranking are determined on the basis of information that is predominantly of a single type" (25). So long as the ranking is primarily determined by teleological or bouletic considerations, for instance, the modal can usefully be considered 'teleological', even if, strictly speaking, the ranking is also influenced by some non-teleological (e.g. moral) considerations. They illustrate the point with the following:

(4) To get to Harlem, you have to take the A train.

[109] See Huitink (2014), among others, for more on various 'modifiers' of modals—as well as Morzycki (forthcoming) on how the category of "modifiers" is something of a grab bag.

The most natural reading of (4) is teleological, one where we hold fixed the relevant circumstances (e.g. we ignore worlds where the transportation system is different) and rank the various possibilities according to their (at least normal) efficiency in achieving the addressee's goal of getting to Harlem. But as Knobe and Szabo point out, even in humdrum cases like this, it's possible for deontic considerations to play a role:

> [S]uppose it turned out that you could get to Harlem very quickly by boarding a different train, taking the conductor hostage at gunpoint, and demanding that the train be rerouted to a different track. There seems to be some way in which this option is ruled out, which is why the sentence can come out true. Now, one might suggest that we ignore this possibility because we assume that this way of getting to Harlem wouldn't fit the agent's larger, unstated goals. But ask yourself, would your answer change at all if it turned out that the agent was a hardened criminal who has no concern whatsoever about taking hostages? If you think that the sentence would remain true even in this case, you think the inner domain here is shaped in this case by deontic considerations over and above circumstantial and teleological ones. (26)

Although Knobe and Szabo are primarily concerned with the role that various deontic considerations might play in the interpretation of modals that are predominantly circumstantial or teleological, they note that once one admits the possibility of impure readings, it's easy to envision a range of different impurities that might in principle arise. This point is worth bearing in mind.

In sum, the standard contextualist account, as presented so far, is fairly straightforward: each modal expression has a common semantic core that constitutes its standing meaning, and the different readings arise from the prejacent being evaluated with respect to different bodies of background information, which can vary from context to context.[110] In this way, all differences between flavors (and

[110] Which in turn might depend on, or at least be constrained by, it's syntactic position—cf. Hacquard (2010).

sub-flavors and sub-sub-flavors…)—other than the ones lexically specified—are attributed to context-dependence, with 'given'-clauses, adverbials, and the like simply being ways of making the value of the relevant parameters (more) explicit. The devil is in the details, of course, and much more needs to be said about the composition of different bodies of background information, or flavors, and how they differ. But we're now in a good enough position to appreciate how modals interact with conditionals.

## 3.3 The restrictor view

The debate over whether the proper interpretation of conditionals like Deductive Coherence is Wide or Narrow—call this the "Scope debate" (as opposed to the scope debate)—concerns the *logical form* of the propositions expressed. Like most recent philosophical discussions of conditionals, the Scope debate has proceeded (for the most part) under the assumption that conditionals are to be formally represented using a two-place conditional operator (=>) that takes a pair of propositions and forms a conditional proposition, similar to how clauses joined by 'and'/'or' can be formally represented using two-place operators (∧/∨) that take a pair of propositions and form a conjunction/disjunction.[111] Call this the "operator view".[112] It's because the operator view takes 'if' to denote a two-place conditional operator that question of relative scope arises whenever there's a co-occurring modal (or other quantificational expression).

Although the operator view is still accepted by many logicians and philosophers, it is widely rejected by linguists. The dominant alternative—commonly known as the "restrictor view"—involves a fundamental re-thinking of the compositional structure of conditionals: rather than denoting a two-place

---

[111] Bennett (2003), for instance, simply *defines* conditionals as any sentence involving a two-place conditional operator.

[112] This is a rough, first-pass characterization of the operator view.

conditional operator, 'if' merely functions as a device for restricting the domains of nearby operators.[113] The restrictor view was first introduced by Lewis (1975) to handle conditionals containing adverbs of quantification ('usually', 'always', etc.), and subsequently generalized to other conditionals by Kratzer (1977, 1981, 2012). The basic idea is that just as in sentences like

(5) *All/Most/Some* men smoke

the common noun ('men') restricts the domain of the quantifier ('all'/'most'/ 'some'), so that it only ranges over (in this case) men, so in conditionals like

(6) If you believe it's going to rain, you *usually/always/sometimes/may/must/ought/are required to* carry an umbrella.

the antecedent ('you believe it's going to rain') restricts the domain of the co-occurring quantifier, which is what adverbs like 'usually'/'always'/'sometimes' and modals like 'may'/'must'/'ought'/ 'required' are standardly analyzed as, per above. As a result, the prejacent ('you carry an umbrella') is only evaluated with respect to the restricted set of possibilities where the antecedent is true (i.e. you believe it's going to rain). Thus a claim of the form 'If you believe it's going to rain, you usually carry an umbrella' will be true (very roughly) just in case most situations in which you believe it's going to rain are situations in which you carry an umbrella. And a claim of the form 'If you believe it's going to rain, you ought to carry an umbrella' will be true (very roughly) just in case the normatively best or "highest ranked" worlds in which you believe it's going to rain are worlds in which you carry an umbrella.

---

[113] At least this is true for the *primary* use of 'if'; like the vast majority of English words, 'if' has more than one use, and these uses differ in syntactically and/or semantically significant ways. For a survey, see Bhatt and Pancheva (2006) and Iatridou (forthcoming), among others.

Although the nature of the quantification involved in sentences like (5) and (6) above differs—quantificational noun phrases ('all men') quantify over individuals of various kinds while quantificational adverbs ('always', 'probably') and modals ('might', 'ought') involve quantification over possibilities of various kinds—they are otherwise similar in structure. In particular, according to the restrictor view, both kinds of sentences should be understood as having something like the following tripartite logical form, where Q is a binary *q*uantifier, R is a (possibly tacit) domain *r*estriction, and P is the matrix *p*redicate (in nominally quantified sentences) or *p*rejacent (in conditionals):

(Q: R) (P)

The traditional view, in contrast, treats both kinds of sentences as having the following bipartite logical form, where Q is an unary quantifier and *conn* is a two-place connective:

Q(R *conn* P)

Kratzer (2012: 106) famously sums up the main lesson as follows:

> The history of the conditional is the story of a syntactic mistake. There is no two-place *if…then* connective in the logical forms for natural languages. *If*-clauses are devices for restricting the domains of operators.

Although it has yet to achieve dominance among philosophers—mostly, though certainly not exclusively, due to unfamiliarity[114]—the restrictor view is a popular and well-established claim in linguistic semantics. It has no serious competitor matching it in simplicity and explanatory power, and the only other contenders are non-unified accounts.[115] This is a large claim, of course, and although I'll indicate some of the reasons why the restrictor view is plausible, an full-fledged defense of it lies beyond the scope of this dissertation.

The main methodological point is this: insofar as we should defer to broad agreement among experts and our best "science", I think we should accept the restrictor view and explore its consequences when it comes to philosophically interesting claims expressed as conditionals. If we do, however, the standard linguistic motivation for taking the scope debate seriously outlined in above is undermined. For it turns out there is no conditional operator—at any level of linguistic analysis—concerning which the question of relative scope makes sense. Thus, insofar as the linguistic argument in favor of wide-scope view presupposes assumes that the Wide scope reading is a linguistically available reading of deontic

---

[114] As Rothschild (forthcoming) notes, despite its linguistic prominence, the restrictor view has "played little role in the philosophical discussion of conditionals. Fairly recent philosophical surveys such as Bennett's (2003) book-length introduction or Edgington's (1995, 2008) review articles do not even mention the restrictor view. Stranger still, in his seminal work on conditionals and probability, Lewis (1976, 1986) does not mention the restrictor view which he pioneered, despite the intimate relation."

[115] Lewis, for example, treated 'if' as three-way ambiguous, defending the restrictor account of adverbially quantified conditionals, a material conditional analysis of bare indicatives, and a variably strict conditional analysis of counterfactuals. Cross-linguistic evidence makes such ambiguity unlikely, however—positing it is an option of last resort. The same moral applies to sophisticated dynamic semantical accounts such as Gillies (2010), which can account for the interaction between conditionals and epistemic modals but not between conditionals and adverbs of quantification, which is what motivated Lewis to introduce the restrictor view in the first place. (See Khoo (2011) for details.)

conditionals like Deductive Coherence (and their instances), it should be rejected. The argument rests on a mistake concerning the structure and meaning of conditionals.[116]

Of course, as I've noted above, not all contributors to the scope debate present it as one concerning the logical form of propositions expressed by natural language conditionals like Deductive Coherence. Instead, many simply help themselves to conditional operator-based formalizations without commenting on what, if any, connection such formalizations have to the intuitive claims we make that motivate theorizing about structural rationality in the first place. But it's hard to make sense of their actual practice of transitioning back and forth between claims expressed as ordinary conditionals and formal representations employing a conditional operator without assuming that there's *some* important relationship between the two. Others have acknowledged the connection is imperfect but nevertheless proceed on the assumption that the simplification involved doesn't matter much for theoretical purposes. In doing so, they are following conventions widely adopted in deontic logic and other areas of philosophy, and so there's a well-established precedent. But the relevance of any given formalization to its closest natural language counterpart is a substantive issue, not one to be decided by deference to tradition, and sometimes simplifications do matter for theoretical purposes. The motivation for taking the scope debate seriously is a case in point.

---

[116] Slightly more carefully, we should distinguish the *semantic* thesis that 'if'-clauses are devices for restricting the domains of various operators from the *syntactic* thesis that there is no two-place conditional operator in the logical forms of natural languages. Taken together they constitute what I'm calling the restrictor view. But even if in practice they tend to go together, in principle they're separable. For it's possible for the semantic thesis to implemented in a variety of ways, including with a two- (or three-) place operator. [Citations and concrete example?] Importantly, however, none of the possible (and plausible) implementations I'm aware of will be of help to the wide-scoper, since they don't allow for *semantically significant* scope distinctions to arise. (That's mainly because if they did they'd give rise to false predictions.) In a nutshell, that's because the operators are invariably defined as ones that *operate on the relevant modal,* and hence don't have the kind of independence from the modal needed in order to enter into scopal relations with it.

I realize that some of those interested in the debate over the nature and structure of rational requirements might be skeptical of the relevance of any investigation into the syntax and semantics of English (or any other natural language for that matter). The primary subject matter, after all, concerns the *requirements of rationality*, and such requirements (if they exist) are plausibly neither mind- nor language-dependent. What's important are facts about *what rationality requires*, not facts about what we *think* rationality requires or facts about the *language* we use to describe what rationality requires. So, from this vantage point, it's hard to see how linguistic data could possibly be relevant to the topic under discussion.

Although I sympathize with this attitude to some extent, and agree that the nature and structure of rational requirements is not at bottom a linguistic issue, the general spirit evinced by this objection is difficult to maintain—at least when it comes to the case at hand. No doubt there are many philosophically interesting questions concerning which linguistic investigation is of limited relevance. But language is, in part, a tool—it's something we use to accomplish certain ends—and like all tools it can be put to better and worse uses. And the uncontroversial fact that language can sometimes be seriously misleading alone justifies paying careful attention to the central bits of language we use in theorizing in order to make sure we're avoiding possible bewitchment. One needn't think paying attention to language will reap substantive philosophical insights all on its own.

It's worth emphasizing at this point that modal expressions and conditional constructions aren't technical terms in philosophy—we can't just *stipulate* what they mean. They're part of natural language, the language we use when engaging in philosophical inquiry. As Jennifer Carr (forthcoming) notes in a related context:

> When we investigate the relations between what we ought to do and considerations like what morality or prudence requires, what we want or intend, what we're able to do, and so on, [as well as what we ought to do *if* various circumstances obtain,] our claims are expressed in natural language. They reflect, even against our will, the logical structure of natural language. And our theorizing often tacitly makes assumptions about the entailment relations between various considerations and [normative facts]. Sometimes these assumptions are false: they involve misunderstandings of what propositions our [normative] claims express and how they logically interact.

Given that the philosophical debate over the nature and structure of rational requirements traffics heavily—and, to a large extent, unavoidably—in 'if's and 'ought's, drawing substantive conclusions from claims made using such expressions, it is *precisely* one of those debates where paying attention to the language we use matters most.

Granted, there definitely *appear* to be ways of explicitly marking a distinction in scope in quantified conditionals, corresponding to a difference in the placement of the quantificational operator in the surface structure of conditional claims. Consider for instance:

(7a) If you take the bet, you will probably lose.

(7b) Probably, if you take the bet, you will lose.

Or:

(8a) If you believe it's going to rain, you should carry an umbrella.

(8b) You should, if you believing it's going to rain, carry an umbrella.[117]

---

[117] Of course, 'if'-clauses can also take final position in a sentence:

You will probably lose if you take the bet.

You should carry an umbrella if you believing it's going to rain.

This doesn't (usually) make a semantic difference either.

But appearances can be deceiving, and as Kratzer (1986, 2012) among others have argued, the surface difference between the (a)-sentences and the (b)-sentences doesn't reflect a deep difference—a difference that is semantically or logically significant.[118]

Of course, many conditionals don't appear to contain a quantificational expression for the 'if'-clause to restrict. Consider:

(9) If Oswald did not kill Kennedy, someone else did.

(10) If the lights are on, John is home.

These "bare" indicative conditionals appear to have binary structure: 'if' connects the antecedent with (what is traditionally thought of as) the consequent.[119] This is what gives the operator view prima facie plausibility.

Here advocates of the restrictor view face a choice—either 'if' is ambiguous between a domain-restricting device in quantified conditionals and a conditional operator in bare conditionals, or else bare conditionals contain an covert (phonologically null) operator for the 'if'-clause to restrict, and hence are not really bare. Although Lewis opts for the former, Kratzer (1986) opts for the latter, conceding that "whenever there is no explicit operator, we have to posit one" (656). Kaufman (2005) goes further,

---

[118] This is slightly overstating the situation, since (as I point out below) many quantified conditionals *are* structurally ambiguous, and different placement of the quantifier can sometimes force, or otherwise favor, one reading at the expense of the other. They're just not ambiguous in the way wide and narrow scopers have thought.

[119] What about "bare" subjunctives—e.g. of the form "If *p* had been the case, then *q* would have been the case"? It turns out there aren't any. For the 'if'-clause restricts 'would', which (as the past tense of 'will'), is standardly analyzed as a necessity modal. Thus, there's no need to posit a covert operator in subjunctives for the *if*-clause to restrict, as there is in the case of indicatives.

arguing that we should think of all tensed clauses as containing (possibly covert) modal operators, and that these are what's targeted by 'if'-clauses. The important point, though, is just that according to the unified restrictor view, "bare" or "unmodalized" indicatives are not in fact bare or unmodalized.

Such a claim might seem implausible or objectionably *ad hoc*. But it's not. The positing of a covert modal is justified by the empirical success of the restrictor view in predicting the right truth conditions for conditionals containing modals, probability operators, and adverbial quantifiers, together with the plausible methodological assumption 'if' is not ambiguous. The reasoning here is fairly straightforward:[120]

1. The restrictor view is empirically well-confirmed across a variety of conditional constructions.
2. The restrictor view requires there to be covert modals in bare indicatives.
3. 'If' is not lexically ambiguous.
4. Therefore, we should posit a covert modal in bare indicatives.

There are are a variety of considerations in favor of the unified restrictor view. I'll mention two. The first is simply that it enjoys presumptive favor. To borrow the legal trope: univocality theories are innocent

[120] Cf. Khoo (2013). Though sympathetic to the restrictor view, Khoo also explores a view that rejects premise three. On this alternative—which he dubs the "strict ambiguity" operator theory—'if' is ambiguous between a restrictor device and a strict epistemic conditional operator. Khoo argues (somewhat tentatively) that the strict ambiguity theory can be developed in such a way as to achieve the same empirical coverage that the unified restrictor view does. If he's right, there may not turn out to be any decisive *empirical* reasons to favor one theory over the other. In that case we must "look elsewhere to tip the scale", by (e.g.) determining whether there is independent evidence for or against the covert modal postulated by the likes of Kratzer. Although the possibility of an empirically adequate ambiguity theory might be thought to give the linguistic argument for wide-scoping new life, it doesn't. For the most plausible way of developing it is one that doesn't allow for the kind of scopal interactions needed—e.g. the conditional operator is defined in such a way as to guarantee that the modal takes "narrow" scope, mimicking the restrictor view's verdict by allowing (or requiring) the modal to be evaluated relative to the "local context" set up by the antecedent.

until proven guilty, while ambiguity theories are guilty until proven innocent (i.e. empirically plausible or necessary).

Of course, ambiguity theorists have a ready reply: the theoretical cost of positing ambiguity needs to be weighed against the benefit of avoiding commitment to a covert operator. And that's right. However, this motivation for positing ambiguity—i.e. avoiding commitment to a covert operator—is substantially weakened once it's recognized that *both* views will ultimately need to appeal to covert operators in order to explain the full range of data. And that's because it turns out that bare indicatives are systematically ambiguous, and it's hard to see how the different readings arise without *something* important happening somewhere below the surface. Consider, for example, the following:

(11) If the lights are on, Billy is home.

This sentence has two readings. Informally, and simplifying slightly: on the first reading, it's a claim that concerns Billy's location at a particular time and it's true just in case Billy is home in all (relevant) situations in which the lights are on. Following Kadmon (1987), I'll call this the "one-case" reading.[121] This reading of (11) is dominant in the context set up by the following exchange:

**A**: Is Billy is home right now?

**B**: I'm not sure. But he always turns the lights off before he leaves. So if the lights are on, Billy is home.

[121] Notice that the one-case reading of (5) is more or less equivalent to (a salient reading) of the following:
If the lights are on, Billy must be home.
This holds in general: adding 'must' (with an epistemic interpretation) to bare indicatives usually results in something (close to) their one-case reading.

But it's also possible to use (11) to express a more general claim, one which is true (again, roughly) just in case Billy is usually or normally home when the lights are on. Call this the "multi-case" reading.[122] It's easily accessible in the following context:

**A**: Billy tries to avoid wasting electricity, doesn't he?

**B**: Yeah. He lives across the street, so I know when he's around—if the lights are on, Billy is home, and if they're off, he's either asleep or away.

What's important for present purposes is that there doesn't appear to be anything in the surface structure of the sentence itself that accounts for the different readings. Plausibly, then, we need to posit semantically significant covert structure—i.e. structure that's not superficially apparent—in order to explain the different readings. One obvious and economical hypothesis is that the different readings are generated by different implicit operators, restricted by the 'if'-clause, and this is just what Kratzer (and others) propose. Though the exact nature and status of the two operators is a matter of debate, there's

---

[122] Notice that while some conditionals readily allow for both readings, such as (5), others favor one reading over (or, at the limit, to the exclusion of) the other. For instance, (A) is most naturally interpreted as a one-case conditional while (B) is most naturally interpreted as a multi-case conditional:

(A) If you took the bet, you lost. [one-case]

(B) If Jane leaves work on time, she has dinner with her family. [multi-case]

Informally, and simplifying slightly: (A) is true on its one-case reading just in case you lost in all (relevant) situations in which you took the bet, and (B) is true on its multi-case reading just in case Jane has dinner with her family in most or all (relevant) situations in which she leaves on time. But in both cases the other reading is available, even if it's difficult to access without contextual clues. For instance:

(A') You used to be a terrible gambler. That's why I kept offering you bets. If you took the bet, you lost. And I, of course, won. [multi-case]

(B') I'm writing a story about Jane, but I haven't decided what to do about the scene where she's facing a deadline at work and might have to stay late. I have only made up my mind about one thing: if Jane leaves work on time, she has dinner with her family. [one-case]

general agreement that what underlies the one-case readings is some sort of necessity operator (*NEC*) and that what underlies the multi-case readings is some sort of generic operator (*GEN*). Together, these two covert operators make possible a unified semantic and syntactic treatment of 'if'. The one-case reading of (11), for example, can be represented as (11') and the multi-case reading as (11''):

(11') *NEC*(the lights are on)(Billy is home)

(11'') *GEN*(the lights are on)(Billy is home)

It's worth emphasizing that the unified restrictor view is more economical than it seems, for it only involves postulating one "new" covert operator rather than two. And that's because the need for (something like) the generic operator is overdetermined: it's widely recognized that a variety of sentences—such as those in the simple present tense—allow for "habitual" readings, and the standard explanation of this appeals to the presence of a implicit generic variable-binding operator.[123] This is precisely the sort of operator the restrictor theorist needs to explain the multi-case reading, together with the assumption that it can be restricted by 'if' when it (covertly) appears in the consequent of

[123] This generic operator is oftentimes glossed as a generic adverbial *quantifier*, but there are some complications concerning its proper semantic classification—complications that I'll be ignoring. In order to remain neutral on in-house disputes, I'm following Leslie (2007, 2008, forthcoming) in taking the relevant generic operator to be a variable-binding operator whose domain is capable of being restricted, but leaving it open whether it should be classified as a genuine quantifier alongside 'all', 'some', 'most', etc. For a defense of a quantificational account of generics, see Nickel (2008, forthcoming) as well as Sterken (2015, forthcoming).

conditionals. The one-case reading is then explained in an analogous fashion by appealing to the presence of an operator that's similar in kind but different in force.[124]

So while everyone should agree that as a methodological matter one should be reluctant to posit covert operators (or covert anything, for that matter), when otherwise recalcitrant data demands it we should oblige. And since *everyone* needs to posit the possible presence of *at least one* covert operator in "bare" conditionals in order to account for multi-case readings, the positing of *another* operator to account for the one-case readings becomes much less of a theoretical cost than it might otherwise be.

Accordingly, the standard development of the restrictor view allows for a certain amount of variation in the kind of implicit operator that might be present in "bare" conditionals, with the evidence suggesting that there are at least two distinct operators that may be present for the 'if'-clause to restrict. One might worry, though, that once we open the door to let in covert operators, there's no non-arbitrary stopping point. Why stop at just two? The short answer: the lack of evidence. Not anything goes, after all—it's standard to posit both *NEC* and *GEN* because that's what the evidence seems to call for, but equally important, that's *all* that the evidence seems to call for. For whatever reason, there seem to be fairly severe constraints on the kind of implicit operator present in conditionals, constraints that appear to only allow for two options: a necessity operator and a generic operator. There's no evidence that 'if'-

---

[124] Again, this conclusion isn't inevitable. For there's an alternative explanation available to those who take 'if' to be ambiguous between a strict epistemic conditional operator ('$if_{=>}$') and a restrictor device ('$if_R$'). (I take this to be the most plausible ambiguity theory on the market.) The one-case reading of bare indicatives is easy: it arises whenever 'if' functions as a strict epistemic conditional operator (i.e. as $if_{=>}$). And in order to account for the multi-case reading the ambiguity theorist can tell the same story as the unified restrictor theorist: it's due to the presence of a covert generic operator, restricted by $if_R$.

clauses ever restrict covert possibility and probability operators, for instance—though of course if such evidence *were* to exist, we should be open to positing them.[125]

In this regard it's worth noting that the positing of covert operators in bare conditionals gives rise to certain predictions. In particular, on the assumption that covert operators can also be present in "quantified" conditionals—understood roughly as any conditional that contains an overt modal or adverbial quantifier as the highest operator in (what is traditionally thought of as) its consequent—two things are to be expected: first, that the systematic ambiguity between one-case and multi-case readings that arises with bare conditionals also arises with quantified conditionals, and second, that there is an additional layer of ambiguity in quantified conditionals, given the possible presence of multiple operators. Both of these predictions are borne out.[126]

[**Note**: prior to the defense I'll be sending an appendix discussing each point in more depth, and how they bear on possible interpretations of conditionals like Deductive Coherence and Strict Means-End. I won't expect the appendix to be read in advance of the defense, however—it'll be optional material.]

---

[125] There may be deeper reasons why *NEC* and *GEN* are privileged in this way and enjoy something like default status. For example, just as Leslie (2007, 2008) and Gelman (2009) suggest that we use *GEN* to express cognitively basic *generalizations*—and that this helps explain various important facts about acquisition and recollection—so we might think *NEC* is used to express cognitively basic *episodic* or *verdictive* judgments about the way the world is. Other forms of generalization (such as universal quantification) and more qualified, graded forms of judgment (such as what's likely to be the case) manifest themselves later developmentally, and so receive phonological expression. This is wildly speculative, though.

[126] See Frank (1997), Geurts (2004), Kratzer (2012), Carr (2014), and Lauer and Condoravdi (2014, ms), among others. Swanson (2010) suggests the phenomenon also arises with counterfactuals. For example, it's plausible to think (18) contains a covert 'would' that scopes above the overt  in the consequent:

(18) If there had been a mural on the floor, every square inch of the mural might have had paint on it.

He notes that this "broadens the base of support" for the view that "conditionals with overt modals in their consequents may also contain covert higher modals that are restricted by the conditional's antecedent" (537).

# CHAPTER 4: AGAINST REQUIREMENTS

## 4.1 The scope debate revisited

As we've seen, even though the restrictor view undermines the standard linguistic argument motivating the debate over the "scope" of rational requirements, and thereby removes a major traditional pillar of support for the wide scope view, it doesn't call into question the existence of rational requirements themselves. It just shows that, contrary to what many have assumed, there's no direct relationship between the requirements themselves (or their instances) and the propositions expressed by sentences like Deductive Coherence and Strict Means End. It's of course fine to appeal to our judgments involving ordinary deontic conditionals for the purpose of motivating our interest in structural rationality itself, as I have done repeatedly. The restrictor view only undermines the appeal to such judgments as a way of motivating the wide scope *theory* of structural rationality.

The obvious thing for the wide scoper to do is concede the linguistic point while insisting that there's a deeper, purely philosophical reason for positing wide scope coherence requirements. For even if the requirements don't have a direct connection to the judgments we canonically express using deontic conditionals, we still need to explain what's bad about, say, believing $p$ and that $p$ entails $q$ without believing $q$ (or while disbelieving $q$) while avoiding the implausible consequence that merely believing $p$ and that $p$ entails $q$ guarantees that you are rationally required to believe $q$ (cf. the "too strong problem" from Chapter 3), and that if you also happen to believe not-$q$ and that not-$q$ entails not-$p$ in such a case you are rationally required to believe not-$q$ as well (cf. the "conflict problem"). And, it might be urged, wide scope requirements expressed (or at least expressible) by the following do that:

(**DC-Wide**) Rationality requires of you that ((you believe $p$ ∧ you believe that $p$ entails $q$) —> that you believe $q$)

while those expressed by the following don't:

(**DC-Narrow**) (You believe $p$ ∧ you believe that $p$ entails $q$) —> rationality requires of you that you believe $q$.

As discussed in Chapter 2, (DC-Wide) merely prohibits a certain combination of attitudes, and is logically equivalent to:

(**DC-Wide'**) Rationality requires of you that you not: believe $p$, believe $p$ entails $q$, and not believe $q$.

This second formulation directly expresses the core commitment of the wide scope view—namely, that structural rationality consists of a set of requirements mandating and/or prohibiting certain *combinations* of attitudes, and that's it. To motivate and state wide scope requirements we needn't appeal to a conditional operator, nor to our ordinary judgments involving iffy oughts.

Things are different with narrow scope requirements. One of the lessons of Chapter 2—though I didn't put it this way at the time—is that narrow scope requirements are best thought of as being *genuinely* and *essentially* conditional in nature. Representations involving the material conditional fail to do justice to this thought—to put it mildly. Take (DC-Narrow) for instance. It it equivalent to

(**DC-Narrow'**) You don't believe $p$ ∨ you don't believe that $p$ entails $q$ ∨ rationality requires of you that you believe $q$.

And this does *not* directly express what I take to be the core commitment of the narrow scope view—namely, that having certain attitudes *makes it the case* that you are required to have certain other attitudes. (DC-Narrow')—and hence (DC-Narrow)—itself says absolutely nothing about the connection between the former attitudes and the law-like requirement to have the latter, conditional on the former. So (DC-Narrow) should be rejected as inadequate, at least in the absence of substantive re-interpretation along the lines suggested in Chapter 2, with the attitudes in the antecedent being stipulatively interpreted as "conditions of application" of the (essentially conditional) requirement to have the attitude specified in the consequent. On its own, (DC-Narrow) is an at best misleading—and at worst positively distorting—representation of the genuinely conditional requirement that narrow-scopers wish to articulate.[127]

## 4.2 Against requirements: stringency

This brings us to the more fundamental problem with the traditional debate over the "scope" of structural requirements, as well with the attempt to explicate the difference between the two dimensions of rationality in terms of requirements (whether of one or both). And that's that requirements (in the "strong" sense) are *at best* only a small part of the story, and arguably not part of the story at all. This becomes clear once we direct our attention beyond a handful of historically distinguished principles involving coarse-grained attitudes, and come to terms with the fact that both

[127] This is one of several reasons why it would be wrong to think of (DC-Narrow) as representing what Broome (1999) calls the "logical factor" of the corresponding narrow scope requirement.

attitudes and the stringency of the normative relations involved come in *degrees*, as well as the fact that higher-order attitudes—just like higher-order evidence—can have "downward" impact.

As warm-up, consider the following claims, where the 'ought' concerns structural rationality:

| | |
|---|---|
| **High Likelihood** | If you believe *p* and that *q* is overwhelmingly likely given *p*, then you rationally ought to believe *q*. |
| **Slack Means-End** | If you intend to φ and believe that ψ-ing is the optimal way to φ, then you rationally ought to intend to ψ. |
| **Higher-Order Coherence** | If you believe that you're probably not in a good position to judge whether *p*, then you rationally ought to suspend judgment with respect to *p*. |

Minor modifications aside, these are *prima facie* plausible in much the same way that Deductive Coherence and Strict Means-End are.[128] And yet the stringency of the normative relation expressed by each is somewhat weaker, as evidenced by the fact that "violations" of the latter seem to be intuitively *worse* (though perhaps not *much* worse) than violations of the former. If the only normative concept we have available to describe structural rationality is that of a strict requirement, however, we will be unable to capture the intuitive difference in strength between them. This suggests that the myopic focus on strict exception-less requirements is misplaced, since it ignores subtleties we're sensitive to.

---

[128] Some might have doubts about the requirements as formulated. Consider, for instance, Slack Means-End. Suppose you think you ought not intend to φ, but intend to φ anyway. And you believe that ψ-ing is the optimal path to φ-ing. In such a scenario, it doesn't seem as though you structurally ought to intend to ψ—you could revise your intention instead. Similarly, High Likelihood seems susceptible to familiar counterexamples involving (e.g.) lotteries. I think such criticisms are correct *when the conditionals are interpreted as expressing strict, exception-less requirements*. But as we'll see, when so interpreted Deductive Coherence and Strict Means-End are also subject to counterexample, and hence remain a par.

This suspicion is reinforced by considering parallel claims involving graded attitudes, such as partial beliefs and partial intentions, which needn't be—and usually aren't—perfectly sharp.[129] Of course, if graded attitudes *were* always perfectly sharp, it might be thought that the connections between them could be framed as strict requirements on precise degrees of partial belief and intention. Most ways of being probabilistically incoherent, for instance, seem pretty bad.[130] One seems to definitely be doing something wrong if one has a high degree of belief in *p* (say, .9) while having a middling degree of belief in not-*p* (say, .4).

But graded attitudes aren't always perfectly sharp. Indeed, it's a familiar point that numerically precise degrees of belief are psychologically unrealistic, and as Joyce (2011) notes, it's "rare, outside casinos, to find opinions that are anywhere near definite or univocal enough to admit of quantification".[131] You might be fairly confident that it will rain tomorrow evening, for instance, and be more confident that it will rain rather than snow, even though there's no particular degree to which you're confident that it will rain (or, for that matter, that it will snow). Perhaps it's because you haven't checked the weather forecast and are merely judging based on your rough recollection of recent weather patterns, or perhaps you have checked the weather forecast but realize that forecasts are imperfect guides to meteorological reality and so remain somewhat unsettled, or whatever. In most situations the information we have at our disposal simply isn't complete, specific, or clear enough to warrant or enable us to invest a precise amount of confidence in a given proposition. These sorts of problems have motivated interest in developing "imprecise" models of partial belief, where such states are represented

---

[129] See Holton (2008, 2009) for the need to recognize partial intentions in addition to full intentions, in much the same way that we need to recognize partial beliefs in addition to full beliefs. See also Shpall (forthcoming).

[130] At least insofar as the probabilistic relations involved are ones the agent is sensitive to.

[131] See, for example, Kyburg (1983), Levi (1985), and Kaplan (1996).

using sets of (e.g.) degrees or probability functions.[132] But even though such models might be *more* psychologically realistic, they still involve a departure from reality, since it's no more plausible to suppose one's level of confidence is sharply bounded than it is to suppose one's level of confidence is sharp. The best way to *model* certain attitudes may very well involve precise tools, but we shouldn't mistake properties of the model for properties of what's being modeled.

The point is that even in the absence of attitudinal precision it's possible for negative evaluation to be called for—it's just a weaker kind of negative evaluation. If you're fairly confident that $p$ and fairly confident that $p$ entails $q$, then it would be somewhat odd if you weren't also fairly confident that $q$, but whatever oddness attaches to such a combination of mental states is less severe than that associated with violations of Deductive Coherence. Similarly, if you partially intend to φ and are pretty confident that φ-ing requires ψ-ing, then it would be somewhat odd if you didn't also partially intend to ψ, but the oddness that attaches to such a combination of mental states is less severe than that associated with violations of Strict Means-End. This isn't to say that the severity of the oddness/badness in cases of unsharp attitudes is *always* weaker than cases of sharp attitudes, nor that it's always weaker in cases of partial attitudes than in cases of full attitudes. Being very confident that $p$ and very confident that not-$p$, for instance, is just as bad whether or not one has numerically precise degrees of belief. The point is just that in *some* cases the severity is weaker, and that we need to be able to account for this difference.

## 4.3 Against requirements: a dilemma

---

[132] See, among others, Joyce (2005, 2011). For dissent, see Carr (ms), among others. It might be objected that although there's no time $t$ and real number $n$ such that you determinately weigh precisely $n$ milligrams at $t$—if only because our bodies lack perfectly precise boundaries—we can still formulate precise laws about weight in those terms. Although that's right as far as it goes, the point remains that we lack the equivalent of bathroom scales for even *approximating* degrees of belief. (Thanks to Stephen Schiffer for discussion.)

The fact that the domain of structural rationality is not exhausted by strict requirements is important, and it motivates the search for an alternative, or at least supplementary, account. However, my goal isn't merely to argue that the requirements-based picture of structural rationality needs supplementation; instead, it's to argue that it needs *replacement*. That is, I want to argue that there's no interesting sense in which the domain of rationality consists *even in part* of strict, exception-less requirements. This includes requirements governing partial attitudes, no matter how sharp those attitudes are.

The challenge can be put in the form of a dilemma: the search for candidate requirements will (almost always) result in either falsehood or triviality. Consider (DC-Wide), for example, or equivalently (DC-Wide'), since it's about as good a candidate as any for expressing a wide scope structural requirement applying to everyone, always, and of necessity. So understood, I think it's false. For suppose you happen to believe that your evidence doesn't support *q*, or that it doesn't support *p*, or that it doesn't support either.[133] Or suppose you happen to believe you're really bad at making deductive inferences, or that you've been slipped a drug that impairs your reasoning ability.[134] Or, more radically still, suppose that as a result of a spurious philosophical argument you come to have non-standard views about logical consequence, and are unsure about whether this is one of those special cases where one proposition logically entails another without guaranteeing its truth.[135] In these cases it's not clear what, if anything, you're required to believe, structurally speaking. But what's important is that it doesn't seem like you're

[133] This is the purely doxastic analogue of having higher-order evidence concerning the quality of one's first-order evidence.

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

[135] That we should be open to the possibility of being mistaken about such matters is one of the upshots of taking content externalism seriously.

*definitely* doing something wrong in virtue of believing *p*, believing *p* entails *q*, and failing to believe *q*—contra (DC-Wide) understood as expressing a strict, exception-less requirement.

Alternatively, suppose that in addition to believing *p* and that *p* entails *q*, you also believe *r* and that *r* entails not-*q*. Or suppose instead you believe *s* and that *s* entails not-*p*. You realize you've made a mistake, but it's not immediately obvious what it is or how to correct for it. What then? Clearly in such cases you're incoherent, and so 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 believing p, believing p entails q, and failing to believe q*. (DC-Wide) misidentifies the locus of incoherence in such cases, since it's more global. It's the full set of conflicting attitudes rather than any particular subset of them.

The defender of (DC-Wide) may respond by pointing out 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-Wide):

{belief that *p*, belief that *p* entails *q*, failing to believe *q*}

{belief that *r*, belief that *r* entails not-*q*, failing to believe not-*q*}

In particular, the defender of (DC-Wide) can point out that no matter which doxastic attitude you adopt vis-a-vis *q*, you're guaranteed to violate the requirement—at least assuming you retain your existing beliefs—and that this fact is "global" in the sense that it concerns the totality of relevant attitudes. Although that's right as far as it goes, it nonetheless misdiagnoses the problem. It's true that you have a set of conflicting attitudes, and that it's in virtue of having those attitudes that you're incoherent, and

hence guilty of structural irrationality. But it *doesn't* seem true that you're guilty of irrationality *twice over*. And yet that's what (DC-Wide) entails, since you violate the relevant requirement twice.

To make the over-counting worry more vivid, notice that the defender of (DC-Wide) will presumably want to capture the fact that there's something distinctively bad about believing *p* and that *p* entails *q* while also believing *r* and that *r* entails not-*q*, over and above the fact that it guarantees that (DC-Wide) is violated. That is, there seems to be something bad about *that very combination of attitudes*, and not just about one or more sub-combinations. So they will presumably want something like the following to express a requirement of rationality as well:

(**DC-wider**) Rationality requires of you that you not: believe *p*, believe *p* entails *q*, believe *r*, and believe *r* entails not-*q*.

But if that's right, then they'll be committed to the view that you are guilty of irrationality not just twice but *three* times over—twice in virtue of violating (DC-Wide) and once in virtue of violating (DC-wider). But that seems wrong. You have definitely gone wrong in virtue of having incoherent attitudes, but you don't seem to have gone wrong in *three* ways. That's over-counting.

The defender of (DC-Wide) might argue that this 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-Wide+**): Rationality requires of you that you not: believe *p*, believe *p* entails *q*, fail to believe *q*, and believe *r*.

Unlike (DC-wider), there's nothing distinctively wrong with the combination of attitudes banned by (DC-Wide+). Although there is *something* wrong with such a combination, the badness is traceable to the first three attitudes alone—the fourth attitude (i.e. the belief that *r*) is extraneous. So unlike (DC-wider), the plausibility of (DC-Wide+) clearly derives from (DC-Wide) .

Call cases involving higher-order attitudes "higher-order cases" and cases involving conflicting attitudes "conflict cases". My claim is that for (almost) any candidate requirement, there will be a range of either higher-order or conflict cases (or both) that constitute counterexamples to that requirement.[136] To take just one more example, consider the following:

(**Non-Contradiction**) Rationality requires of you that you not believe: *p* and not-*p*.

As with (DC-Wide), there are grounds for rejecting (Non-Contradiction) as a strict requirement, despite its initial plausibility. Perhaps the most promising (though still controversial) counterexamples involve so-called "glut theorists", who think there are true contradictions.[137] As J.C. Beall (2013) 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. (3)

---

[136] Although I've been ignoring the possibility of there being principles of permission in addition to principles of obligation, it should be clear that the same moral will apply to principles of permission, at least insofar as they purport to be exception-less.

[137] For defenses of "glut theory", see Jc Beall (2009) and Graham Priest (1979, 2006), among others.

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 (Non-Contradiction) 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.[138]

It's worth reiterating that I'm *not* denying that there are necessary conditions for being fully or ideally rational, and in that sense some "requirements" of rationality. Perhaps anyone who violates (DC-Wide) or (Non-Contradiction) is guaranteed to fall short of some rational ideal or other. Even so, as emphasized earlier, that's not what the debate over the nature and scope of rational requirements has been about. Instead, the relevant requirements are supposed to be strict rules or principles which not only ensure that a subject is irrational whenever they happen to be violated, but also that the subject is irrational *in virtue of* violating them. They're supposed to be telling us something about the nature of structural rationality, rather than merely providing 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 primarily intended to cast doubt upon. If such cases reveal it's possible to violate (e.g.) (DC-Wide) or (Non-Contradiction) without doing *anything* wrong, structurally speaking, so much the better. (DC-Wide) and (Non-Contradiction) are two of the most plausible requirements, and so we might expect other candidates to share the same fate—that is, we might expect that, for any candidate requirement *R* (in the strong sense), there will be cases in which a subject violates *R* without being structurally irrational. But all I need is for higher-order and/or conflict

[138] For an even fancier counterexample to probabilism—i.e. the claim that a rational agent's credences should always be probabilistically coherent—see Caie (2013). For more traditional (and not unpersuasive) objections, see Foley (1993), among others.

cases to make plausible the weaker claim that for any given candidate requirement *R*, it's possible to violate *R* without being irrational *in virtue of* violating *R*.[139]

## 4.4 The explanatory vacuity of requirements

Now, it might be thought the above dialectic just shows that further conditions—such as the absence of higher-order funny business, the absence of fancy philosophical views, and the absence of other conflicting beliefs—need to be built into the relevant requirements. But then the structural requirements threaten to be no more interesting or informative than the so-called "requirements" of reasons rationality considered in Chapter 3. Recall:

**Conclusive Reason** If you have conclusive reason to believe *p* and that *p* entails *q*, then you rationally ought to believe *q*.

**Decisive Reason** If you have decisive reason to φ, then you rationally ought to (intend to) φ.

Although as stated these are both plausibly *true*, the principles expressed aren't explanatory in the way that requirements in the strong sense are supposed to be. Instead, much (if not all) of the normatively interesting action occurs "off-stage" in determining whether you in fact have conclusive or decisive reason, and thereby satisfy the antecedents. For it's very plausible to take facts about whether you have

[139] This way of putting the point is a little tricky, since for any necessary condition *C* of status *S*, if *X* fails to meet *C* then at least *some* sense can be made of the claim that *X* fails to have *S* at least partly "in virtue of" failing to meet *C*. Failing to meet *C*, after all, guarantees that *X* fails to have *S*, and so might be included in a maximally complete explanation of why *X* fails to have *S*. But this just shows that 'in virtue of' has a stronger and weaker sense, just as 'requires' does, and in both cases it is the stronger sense that is of central concern.

conclusive or decisive reason to be *resultant* normative facts that depend on more normatively fundamental facts concerning all the particular reasons you have and how those reasons "add up". Suppose, for instance, that you have decisive reason to (intend to) eat healthily. That's not a basic normative fact. Instead, it's something we expect to have an underlying explanation. In particular, we expect it to be explained by (a) the reasons you have to eat healthily (e.g. it will prolong your life, cause you to lose weight, etc.), (b) the reasons you have to *not* eat healthily (e.g. it would save time and money, etc.), and (c) the "strength" or "weight" of the reasons for/against, both individually and collectively. Not everyone has decisive reason to eat healthily, after all—if you can't afford or lack access to healthy food, or if you only have two days left to live, then you don't have decisive reason to eat healthily. But assuming you *do* have adequate income, access, and so on, you do have decisive reason. If you then fail to eat healthily, you'll be rationally criticizable—you won't have responded correctly to the reasons you have.

This illustrates what I mean when I say that much, if not all, of the normatively interesting action occurs "off-stage" when it comes to Decisive Reason. And the same applies to Conclusive Reason. Such principles simply assume the relevant normative action has been resolved in one way or another without *themselves* playing any role in explaining why it has been resolved that way. Nor do they ultimately explain why agents are guilty of irrationality when they are "violated". When agents are irrational in virtue of being unreasonable the fundamental explanation is "bottom up" rather than "top down": they are irrational because they fail to respond correctly to the reasons they have—reasons which may not be shared by other people in other circumstances—and *not* (in the first instance) because they fail to satisfy some strict law-like requirement. Explanatorily, such principles are epiphenomenal; they summarize the important normative facts rather than dictate them.

My claim is that an analogous result awaits efforts to weaken the principle expressed by (DC-Wide), as well as other candidates for structural requirements. Without something like a ceteris paribus (or "other things being equal") clause, the continual threat of counterexample will result in considerable (and perhaps never-ending) complexity, with the result looking increasingly ad hoc and less explanatory at each step. The recipe for generating counterexamples is straightforward. For every level *n* or kind *K* of judgment, it's possible to adopt a range of higher-order attitudes concerning, for instance, the (un)reliability of one's judgments at level *n* or of kind *K,* and for any rationally evaluable attitude *A* (belief, intention, preference, and so on) 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.

Rather than adding conditions to handle each potential counterexample, then, one might switch tactics and try to guard against potentially disruptive attitudes by adding (something like) a ceteris paribus clause. Suppose, for example, we modify (DC-Wide) as follows:

(**DC-Wide$_{CP}$**) Rationality requires of you, ceteris paribus, that you not: believe *p,* believe *p* entails *q,* and not believe *q*.

Although the addition of the ceteris paribus clause prevents (DC-Wide$_{CP}$) from being falsified by cases involving normatively disruptive action—since those are presumably cases in which other things are *not* equal—it threatens to rob (DC-Wide$_{CP}$) of its potential to do real explanatory work in cases where other things *are* equal. For we are going to want a story about why rationality requires you to not believe *p*, believe *p* entails *q*, and fail to believe *q*, when it does, given that it doesn't always do so. (DC-Wide$_{CP}$)

itself doesn't provide an answer—all it does is state *that* you're prohibited from having such a combination of attitudes, when other things are equal. It doesn't tell us *why*. And the worry is simply that whatever explains why the requirement expressed by (DC-Wide$_{CP}$) applies, when it does, will also be what explains why, in such cases, it would be irrational to believe *p*, believe *p* entails *q*, and fail to believe *q*.

Despite being true, then, the worry is that (DC-Wide$_{CP}$) will be bereft of any meaningful explanatory role. This is an instance of a more general worry concerning hedged normative principles (e.g. 'Ceteris paribus, lying is wrong', 'Ceteris paribus, winning is good')—namely, that whatever explains why the hedged principles apply when they do and why they don't when they don't will *also* be what explains why the normative facts the principles are concerned with (e.g. lying is wrong, winning is good) obtain when they do and when they don't when they don't. It's therefore natural to suspect hedged normative principles *themselves* of being explanatorily idle, in much the same way that run-of-the-mill generics such as 'Babies are cute' or 'Ducks lay eggs' are. They're true, but it's what *makes* them true that's of greater theoretical or explanatory importance.

To make the point in a slightly different way, suppose you believe *p*, believe *p* entails *q*, and fail to believe *q*. And suppose further there's no normatively disruptive action taking place—i.e. you lack higher-order and conflicting beliefs. Given (DC-Wide$_{CP}$), it follows that you're irrational. But entailment is not the same as explanation. What we want to know is whether you're irrational in such a case *because* you've violated the (hedged) requirement expressed by (DC-Wide$_{CP}$). And the answer, it seems, is no. For it's only possible to "violate" (DC-Wide$_{CP}$) when there's no normatively disruptive action taking place. (When the ceteris paribus clause isn't satisfied, the requirement expressed by (DC-Wide$_{CP}$)

doesn't apply—it's simply silent.[140]) The absence of normatively disruptive action is thus something that the requirement expressed by (DC-Wide$_{CP}$)—when it applies—assumes rather than explains in much the same way that the principle expressed (e.g.) Decisive Reason—when it applies—assumes without explaining the fact that you lack sufficient reason not to φ.[141] In both cases most of the normatively interesting action occurs off-stage in determining whether the principles "apply" at all. In the case of Decisive Reason, it's in determining whether the antecedent is satisfied—i.e. whether you have decisive reason—and in the case of (DC-Wide$_{CP}$) it's in determining whether the ceteris paribus clause is satisfied—i.e. whether you have any normatively disruptive attitudes. And so just as we naturally expect a "common cause" explanation in the former case, with the fact that you have decisive reason to φ and the fact that you lack sufficient reason not to φ being jointly explained by the same underlying facts concerning the totality of your reasons (and their relative strengths, etc.), so we should expect a common cause explanation in the latter case, with the fact that it's irrational to believe *p*, believe *p* entails *q*, and fail to believe *q* and the fact that there's no normatively disruptive action taking place that makes such a combination permissible being jointly explained by the same underlying facts concerning the totality of your attitudes.

The worry generalizes beyond the likes of (DC-Wide$_{CP}$) to include many, if not all, hedged requirements. The challenge facing the defender of hedged yet explanatory normative principles is to provide a plausible account of the role ceteris paribus clauses play in such principles that doesn't jeopardize their explanatory status. To be clear: I'm happy to grant that there are non-trivial ceteris

---

[140] Compare: when you don't have decisive reason to φ you don't "violate" Decisive Reason, but that's just because it is silent about such cases.

[141] The absence of normatively disruptive action also admits of varying explanations—it might be the result of there not being any disruptive attitudes, or instead of all potentially disruptive attitudes themselves being disrupted (and so not being *actually* disruptive).

paribus generalizations elsewhere, including in the special sciences. This may seem to give defenders of (DC-Wide) some hope, since they may propose to understand it as a ceteris paribus generalization along similar lines. Unfortunately, however, none of the accounts I'm aware of that are able to rescue generalizations from triviality elsewhere achieve the same result with (DC-Wide). To take just one example: Strevens (2012) provides a promising account of the role played by ceteris paribus hedges when they are added to empirical generalizations intended to articulate the consequences of a causal mechanism. These generalizations include claims like 'Ceteris paribus, ravens are black', understood as a claim concerning the effects of the natural raven-coloration mechanism. According to Strevens, the function of such a hedge is to restrict the scope of the empirical generalization to those cases where "nothing undermines, interferes with, or undoes the effect" of the target mechanism (652).[142] By focusing on the consequences of the mechanism in question, however, the generalization remains open to disconfirmation. Although any state of affairs *not* caused by the mechanism in question is rendered irrelevant, those that *are* caused by the mechanism are not, and hence are able to (dis)confirm the generalization.[143]

The appeal to an underlying mechanism is therefore crucial to circumventing the worry that the hedge trivializes the generalization, making it little more informative than 'Ravens are black—except when they're not'. The problem for the defender of (DC-Wide), however, is that there's nothing analogous to a causal mechanism to ward off the threat of trivialization. So another story needs to be

---

[142] This is an instance of what Stevens calls the "narrowing" approach to understanding the significance of *ceteris paribus* hedges. Although he considers two other approaches (the "softening" approach and the "annotating" approach), neither of them would be of use to those wishing to defend (DC-Wide) as being genuinely explanatory. Instead, they would reduce (DC-Wide) to being a true exception-permitting generalization akin to the claim that beds are rectangular.

[143] Cf. Strevens (2012: 665).

told about the role of the ceteris paribus hedge in such cases—a story that doesn't rob (DC-Wide) of its explanatory status. The hedge can't just be restricting the scope of the principles to those cases which don't constitute counterexamples to their unrestricted versions. For then the principles would be trivial, akin to claims of the form 'Rationality prohibits you from having attitudes $\{A_1,\ldots,A_n\}$—except when it doesn't'. What we want is a story about *why* you're rationally prohibited from having certain attitudes when you are, given that you're not always so. It can't just be a brute fact. In the absence of such a story, we should be open to the possibility that the "requirements" of structural rationality, just like the "requirements" of reasons rationality, are susceptible to a deeper, underlying explanation, and aren't themselves explanatory—and hence not really requirements (in the strong sense). This is a possibility I'll return to—and defend—when presenting an alternative account of structural rationality in the next chapter.

## 4.5 Two modest lessons

I realize, however, that to actually *establish* the claim that there are no rational requirements in the strong sense would require considering a much wider range of candidate requirements, both practical and theoretical, hedged and unhedged. And it would require not only offering further counterexamples, but also a more thorough consideration of possible revisions, a more careful investigation of the role ceteris paribus clauses might play, and the rebutting of various objections to my objections. My discussion of the principles above merely represents the beginning of such an attempt—one that I take to be highly suggestive, but hardly conclusive.

So suppose I'm wrong—suppose it turns out there are a range of intuitively plausible requirements that appear immune from counterexample. Two more modest points remain. The first is

that even if a general principle is free from counterexample, it doesn't follow that it's a *requirement* in the strong sense. For the generalization may be susceptible of a "bottom up" explanation in terms of something more fundamental, in which case both its truth and its explanatory relevance will be derivative. The principles expressed by (DC-Wide+) and (Conclusive Reason)/(Decisive Reason) should be relatively uncontroversial examples of derivative normative principles, even though the ways in which they're derivative differ. Before concluding that some apparently explanatory principle is in fact a genuine requirement, then, one needs to rule out the possibility of it being amenable to a similar treatment.

The second point is just that appeals to requirements fail to exhaust our intuitions about structural rationality, and so it's a mistake to think structural rationality can be understood entirely in terms of them. This is primarily what motivates the development of the alternative, pressure-based view in what follows. Once the view is on the table, however, I'll return to the status of principles like those expressed by (DC-Wide), explaining how the pressure-based view can not only account for the plausibility of such principles without reifying them as requirements, but also explain when and why apparent counterexamples arise. Assuming there is no explanatory work left undone, it's with this fact that I'll rest my case that structural rationality should not be understood even in part in terms of requirements.

# CHAPTER 5: RATIONALITY AND PRESSURE

## 5.1 Threshold-y vs. graded normative notions

In this chapter I argue that the nature of structural rationality is importantly analogous to the "force-like" or "pressure-based" nature of reasons rationality—and importantly *disanalogous* to the "rule-based" nature of the law, contra (e.g.) Broome and Schroeder. I say 'rule-based' rather than 'requirement-based' because requirements are only one type of rule, or principle; the law also includes various principles of permission. Broome (2013) claims that rationally does as well. Although requirements and permissions differ in normative strength, they are alike in being essentially "threshold-y" or "all-or-nothing" statuses. For any relevant normatively evaluable entity *X*—whether it be an act(-type), state of affairs, rule, combination of attitudes, or some other "evaluative focal point" in Kagan (2000)'s terminology—and rule-based normative domain *D*—such as the law or etiquette—either *X* is required/permitted by *D* or it is not.[144] Requirements and permissions are thus importantly unlike "graded" or "quantitative" normative notions such as value, justification, and reason, all of which come in (possibly vague or indeterminate) degrees. So whereas it doesn't make sense to say of some action-type or state of affairs *A* that it is more required/permitted than B, it *does* make sense to say that *A* is more valuable/justified/well-supported than *B*, or that you have more reason/justification to *A* than to *B*. Of course, some rules may be more important, or "ranked higher", than others, and hence take precedent in cases of conflict. But to say that rules admit of hierarchical relationships, such as rankings, is not to say that rules themselves come in degrees. Rank-ability is not gradability, in the relevant sense.

[144] I'm intentionally ignoring the possibility of various forms and sources of vagueness or indeterminacy.

The alternative picture of structural rationality, or coherence, that I'll be suggesting is not new. Pryor (2004), for instance, has suggested something similar when it comes to epistemic or theoretical rationality. It has nonetheless been generally overlooked, and the debate over the nature and structure of structural rationality—both practical and theoretical—has proceeded on the largely unquestioned assumption that it bottoms out (at least in part) in requirements, rather than something more graded. Nor do I think the pressure-based view is particularly revolutionary. If anything, it's the one that deserves to be treated as default, rather than the requirements-based view. For as Kolodny (2008: 437) himself notes 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.

Although Kolodny makes similarly suggestive remarks elsewhere—including in his seminal (2005)—he nonetheless never seems to recognize how uneasily such remarks fit with a requirements-based picture. It's a testament to the influence of Broome (1999, 2004, 2013), among others, that almost nobody has seriously questioned appropriateness of thinking of structural rationality in terms of requirements, rather than something else—something more graded and force-like. The ideology of requirements has thus reigned largely unchallenged, despite the intuitive force of observations like Kolodny's above.

## 5.2 Attitudinal vs. justificatory pressure

To return to one of the main questions of this dissertation: how, exactly, we should understand the difference between the different dimensions of rationality, if it's not a matter of there being differences in the structure of their requirements, or of one being requirements-based and the other not?

The answer to this question, I suggest, comes into focus once we realize that the phenomenon is present even—indeed, usually—in the absence of a corresponding 'ought'- or 'requires'-judgment. 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's clearly *some* rational pressure on the juror to believe that 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 as well as the prosecutor are thoroughly corrupt, and so that the defendant has probably been framed.

Or instead imagine you're a scientist who is convinced the available empirical evidence strongly supports a certain theory *T*. But suppose you also happen to believe that the results of a major experiment are about to be announced, and that the results (whatever they are) will either provide near-decisive confirmation of the theory or else near-decisive refutation. In such a case although there clearly seems to be *some* rational pressure to believe that *T* is true, and that the results of the experiment will confirm it, there there's no need to have reached a settled opinion on either matter (especially the latter), and so it's not the case that, structurally speaking, you ought or are required to believe either. Again, this is even clearer if you have beliefs that generate conflicting pressure—perhaps, for example, you also happen to be deeply religious, and believe that *T* is incompatible with divine revelation.

Analogous cases can be constructed in the practical realm concerning the relationship between what you believe it would be good (to some degree) to do and what it would be structurally rational for

you to (intend to) do. In general, then, claims like the following (or suitably Chisholmed versions of them) seem plausible:

(**Evidence**) If you believe there's (at least some) good reason to (not) believe *p*, then there's (at least some) rational pressure to (not) believe *p*.

(**Goodness**) If you believe that it would be good (in some way to some degree) for you to (not) φ, then there's rational pressure (to some degree) to intend to (not) φ.

Neither (Evidence) nor (Goodness) specifies *how much* pressure there is in a given case, but plausibly it's 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. The amount/strength of pressure present will therefore vary, and won't always (or even usually) be strong enough to *require* anything of you, or even single out a particular attitude as best supported. This, I take it, reinforces the point made in the last chapter—namely, that the fact that the stringency of the normative relations associated with structural rationality comes in degrees shows that facts about what structural rationality "requires" capture, at best, only part of the story.

Somewhat surprisingly, one of the most prominent defenders of a requirement-based conception of structural rationality, John Broome, has expressed sympathy with something like this point in the past. In his early and seminal work on normative 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 on requirements.[145] 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 I highlight below.[146]

Accordingly, rather than being seen as corresponding to a distinctive set of rational *requirements*, I think each dimension should be understood as associated with a distinct kind of pro tanto rational *pressure* or *force*—reasons rationality with (what I'll call) *justificatory pressure* and structural rationality with *attitudinal pressure*. On this view, believing $p$ and that $p$ entails $q$ generates substantial attitudinal pressure to also believe $q$, just as having good reason to believe $p$ and that $p$ entails $q$ 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 attitudes.

Thus, despite being distinct and autonomous, the two dimensions of rationality display important similarities—similarities which I'll turn to below. First, though, it's worth getting clear on the

---

[145] As noted in Chapter 1, Broome (2013) does argue in favor of there being "basing permissions", understood as the negation of "basing prohibitions". But these are *diachronic* principles connecting your attitudes at one time with your attitudes at another time, and hence not of present concern.

[146] This problem is analogous (though not identical) to that facing those who think 'ought' is ambiguous between, say, a "subjective" and "objective" sense, or between a "moral" and "epistemic" sense. For to do justice to our various intuitions we'll have to posit an "annoying profusion of 'ought's" (as Jackson (1998) put it), and we should be antecedently suspicious of such large-scale lexical ambiguity. As explained in Chapter 3, it's far more plausible to take 'ought' to be context-sensitive, with the particular "flavor" (e.g. moral, teleological, epistemic) of modality expressed being determined by context.

sense in which they're distinct and autonomous. I take the two dimensions to be *distinct* in the sense that you can satisfy the demands of one without satisfying those of the other. Most obviously, you can be structurally rational without being reasons rational, as evidenced by paranoids and conspiracy theorists who might be perfectly coherent yet thoroughly unreasonable. Less obviously, you can be reasons rational without being structurally rational, as evidenced by the wide range of cases considered in Chapter 1 that (taken together) pose a seemingly insurmountable problem for "structural deniers".

I take the dimensions to be *autonomous* in the sense that the presence of one kind of pressure in one direction doesn't—by itself—generate, constitute, or otherwise guarantee the presence of the other kind of pressure in that direction. Although there might be circumstances that guarantee the presence of both kinds of pressure, on the assumption that *mere* attitudes do not (in general) generate justification—unlike, say, experiences, which do—the explanation of why one kind of pressure exists will be distinct from that of the other.[147] I also take the two dimensions 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 which you recognize to be incompatible with *p*. Your mere *belief* doesn't, by itself, make any difference to what your evidence supports. Although I don't have much of an independent argument for such claims, I take them to be plausible enough to adopt as working hypotheses. Thus, although there are interactions between pressures of the *same* kind, some of which

[147] There may be exceptions. For instance, while mere beliefs do not generate justification—other than perhaps justification to believe you have the belief itself (cf. Bryne (2005))—desires might generate (defeasible) justification to do what you desire to do. There may also be cases in which the dimensions aren't cleanly separable, especially if some form of "impure" or "non-doxastic" coherentism about justification (cf. Berker (forthcoming)) or Wright's (2004) theory of entitlement are true.

will be discussed below, I'll be assuming there aren't any interactions between pressures of *different* kinds.[148] All interactions are intra-dimensional, not inter-dimensional.

In the remainder of this chapter I'll explore some of the most salient similarities and differences between the dimensions, and along the way explain how the pressure-based conception captures what wide- and narrow-scopers get right while avoiding what they get wrong.

## 5.3 Similarities between dimensions

The first point of similarity is the obvious one: each dimension corresponds to a distinct kind of rational pressure, or force. Notice that we're used to thinking of one's reason(s)/justification/evidence as being force-like, having something like *magnitude* (strength), which can vary, and *direction* (for/against). That much is old news. What I'm suggesting is that we need to recognize a distinct kind of pro tanto rational pressure that behaves similarly but is importantly unlike one's reason(s)/justification/evidence insofar as it is generated by the mere presence of attitudes, regardless of their justificatory standing. It too is force-like, having 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,

[148] It may be that we have reason to be structurally rational (at least in general), but that's not in tension with the claim I'm making here. That is, there may be justificatory pressure to be in a state of compliance with (resultant) attitudinal pressure, but we shouldn't confuse *being in a state of compliance with attitudinal pressure* with attitudinal pressure itself.

eliminating—other justificatory/attitudinal pressures.[149] 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 it ends up that way will differ. Taking the force metaphor seriously allows us to draw a variety of other distinctions—including between importantly different forms of undermining—but since these are in-house disputes, the details needn't detain us.[150]

---

[149] Cf. Pryor (2004), which—as noted above—is an important precursor to the view on offer. Although Pryor focuses on the rational relevance of bad (i.e. unjustified) doxastic attitudes, much of what he says can be generalized to non-doxastic attitudes, as I have done here.

[150] See, for example, the distinctions in Pryor (2013) between different forms of undermining.

A second similarity between dimensions 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, holding everything else fixed.[151] This thesis is admittedly rough and subject to indeterminacy.[152] Nonetheless, something like it seems true in a wide range of cases. Return to Tom, for instance. How much attitudinal pressure is there for him to believe that he can fly, given that he believes that he's Superman and that Superman can fly? 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?[153] The rough 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, or partial belief, is of course complicated, but it's generally agreed that one can believe $p$ and believe $q$ while nonetheless being more confident in $p$ than $q$. And, 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 there will be.

---

[151] I realize this is a counterfactual with an impossible antecedent, given that we're supposed to be keeping everything other than the justificatory status of Tom's beliefs fixed and this involves violating the supervenience of the normative on the non-normative. I take solace in the fact that we make make judgments involving counterpossibles all the time, at least in philosophy, and have fairly robust intuitions concerning their (non-vacuous) truth values.

[152] Cf. Pryor (2004), who puts forward (what I call) the proportionality thesis as a heuristic.

[153] The supposition that Tom's attitudes are justified is just that: the supposition *that the attitudes are justified*. It's not the supposition that the attitudes are justified in any particular way, although of course were the attitudes to actually be justified they would be justified in some particular way—i.e. there would be some explanation of why the attitudes were justified.

Attitudinal pressure can thus be usefully understood as functioning like a kind of "hypothetical" justificatory pressure.[154]

It's worth emphasizing that although attitudinal pressure is generated by or grounded in facts about our psychology, it's not itself psychological. In general, we shouldn't confuse normative facts and relations with what grounds them. Moreover, 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 would make a difference, however, since they (perhaps together with other attitudes) would generate pressure of their own that can then interact with other attitudinal pressures. Suppose, for instance, that you believe *p*, and that *p* entails *q*, but come to think that *q* isn't well-supported by the evidence. In such a case the former attitudes, taken together, generate substantial rational pressure to believe *q*, while the latter attitude generates rational pressure to *not* believe *q*, and hence together with the belief that *p* entails *q* generates pressure to not believe *p*. How the competing pressures are resolved in such a case plausibly depends on their respective strengths, and it may turn out that they fail to single out a particular attitude for revision. Instead, they may deliver an irreducibly disjunctive verdict concerning what the rationally required response is—a point I'll return to below.

Note that because attitudinal pressure is always grounded in the presence of various attitudes, the *failure* to have a certain attitude is never itself a source of attitudinal pressure. So whereas believing *p* and that *p* entails *q* generates attitudinal pressure to also believe *q*, merely failing to believe *q* while believing that *p* entails *q* doesn't generate attitudinal pressure to not believe *p*. Suppose, then, that you believe *p* and believe that *p* entails *q*, but simply haven't put "two and two together" and as a result

[154] Cf. Pryor (ms), though this use of 'hypothetical' should not be taken literally since attitudinal pressure is just as real or substantial (and in that sense *non*-hypothetical) as justificatory pressure is.

haven't yet seriously considered the question of whether *q*. In such a case, although there's (unrecognized) attitudinal pressure to believe *q*, I don't think there's any attitudinal pressure to fail to believe *p*. It's only if one were to, say, suspend judgment with respect to *q*, or view *q* as unlikely, or otherwise harbor doubts about *q* that such attitudinal pressure would arise.[155]

A third important similarity between dimensions lies in the fact that just as it's plausible to take facts about what you're justified in believing and doing to be determined by facts about justificatory pressure—i.e. about what evidence or reason(s) you have—rather than the other way around, so it's plausible to take facts about what you're structurally committed to believe and do to be determined by facts about attitudinal pressure—i.e. about what attitudes you have—rather than the other way around. So although I'm *not* skeptical about there being facts about what "rationality requires" in particular circumstances, where this is understood as all-pressures(-of-a-certain-type)-considered verdicts about which attitudes one rationally ought to have or lack in those circumstances, I *am* skeptical that there are any principles akin to those expressed by (DC-Wide), (ME-Wide), (AA-Wide), and the like that do real explanatory work. For even if it turns out there are some plausible counterexample-free principles, which is a possibility left open at the end of the previous chapter, it's not clear why we should take the relevant principles to be genuinely explanatory—or requirements in the strong sense—as opposed to merely being true generalizations underwritten by stable facts about the attitudinal pressures generated by the relevant attitudes and how they interact.[156]

Consider, for instance, (DC-Wide). How might we explain its plausibility by appealing to facts about pressure? That's easy: the plausibility of (DC-Wide) is explained by the fact that cases in which

---

[155] See Friedman (2013a,b) for a defense of suspended judgment as a positive, irreducible doxastic attitude.

[156] By granting the existence of at least some finite, substantive, and exception-less generalizations, the resulting view would be a form of normative "regionalism", rather than "holism".

one believes *p*, believes *p* entails *q*, and yet doesn't believe *q* are normally going to be cases in which one flouts some substantial (and undefeated) attitudinal pressure, and hence cases in which one fails to be fully structurally rational. And that's because believing *p* and believing *p* entails *q* generates substantial rational pressure to believe *q*—pressure which (if undefeated) is flouted by failing to believe *q*. A similar explanation is available of principles banning inconsistent attitudes.[157] The simplest is that expressed by (Non-Contradiction), requiring that one not believe propositions of the form *p* and not-*p*. The reason why having such a belief is (nearly always) bad is that it requires believing both *p* and believing not-*p*, and each taken on its own, and given their manifest incompatibility this generates substantial rational pressure to revise the other.[158] So when one adopts both simultaneously one is guaranteeing the presence of such pressure and also the flouting of it—at least under normal conditions. It takes (very) fancy philosophical views to reduce or disrupt this inter-attitudinal tension.

There's also a fairly straightforward explanation of why natural language conditionals like Deductive Coherence seem true. It's because they *are* true, at least on one natural reading. This is the reading on which they're understood as expressing general claims about what (in normal cases) is *best*, structurally speaking, given the truth of the antecedent. Holding fixed the fact that you believe *p* and that *p* entails *q*, for example, the structurally best response—in the normal case, and hence generally—is to believe *q*. After all, the other options are to either believe not-*q* or withhold judgment with respect to *q*,

---

[157] As noted in Chapter 2, one might not think there is anything wrong with inconsistency *as such*—perhaps it's only *believed* or *sufficiently obvious* (or …) inconsistency that's bad, as both Field (2009) and Harman (2009) suggest, despite their differences.

[158] There may be cases in which their incompatibility is *not* manifest, if (for example) we were to suppose that 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 plausibly express a contradiction, and it might be one that it is rational to believe insofar as one fails to realize Superman is Clark Kent. I don't wish to take a stand on the metaphysics of propositions, however, so I'm ignoring these possible complications.

both of which would normally result in incoherence. So the counterexamples above do not threaten Deductive Coherence on this reading, since they concern "abnormal" cases. The counterexamples only challenge the particular *interpretations* of Deductive Coherence that philosophers attracted to a requirements-based conception of rationality try to foist on them, or wrongly take them to be evidence of. [**Note**: This is a claim I'll defend more fully and carefully in a future addition or appendix to this chapter.]

Explanations along these lines generalize, and this serves to underscore the point that the existence of seemingly substantial and exception-less principles—as well as plausible natural language conditionals that seem to express or support such principles—isn't enough to motivate the positing of requirements in the strong sense. Because a pressure-based view can provide an alternative explanation structural of such principles and conditionals, a further argument in favor of positing law-like requirements is needed. Absent such considerations, while one may admit there are structural requirements in the *weak* sense, one needn't admit that there are any in the strong sense. Instead, we can view candidate "requirements" of structural rationality as analogous to "requirements" of reasons rationality: rather than being reified as requirements, such principles can be understood as true generalizations underwritten by more fundamental facts concerning various rational pressures or forces and how they interact.

This is especially attractive if, as I think, the relevant structural generalizations are in fact exception-permitting. Whereas exceptions function as counterexamples on the requirements-based picture, they are susceptible to explanation on the pressure-based picture. For in the "higher-order cases" the higher-order attitudes are just what they seem to be—namely, sources of downward attitudinal pressure that interacts with the various pressures generated by the lower-order attitudes. Likewise, in the "conflict cases" the conflicting attitudes are sources of attitudinal pressure that conflicts

with the various pressures generated by the other relevant attitudes. When you believe *p* and that *p* entails *q*, for instance, you are thereby rationally pressured to believe *q*, and in the normal case (i.e. absent higher-order or conflicting attitudes) this pressure will be enough to structurally "require" you to believe *q*. But in abnormal cases, this 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.

On the resulting view, then, facts about rational pressure are taken to be normatively fundamental in both dimensions, and facts about threshold-y (rather than graded) normative statuses like being required or adequately justified are taken to be resultant, derivative facts. The account of structural rationality, like that of reasons rationality, is pleasingly economical: all facts about rationality are explained in terms of two flavors (attitudinal and justificatory) 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. If we can reduce the number of distinct normative phenomena without sacrificing our explanatory ambitions, we should.

## 5.4 Explaining intuitions of asymmetry

It's worth noting that the fact that *failing* to have a certain attitude is never itself a source of attitudinal pressure together with the fact that the strength and direction of attitudinal pressure "mirrors" the strength and direction of (actual or counterfactual) justificatory pressure provides a compelling resolution of one of the main points of contention in the scope debate.

As noted in Chapter 2, all wide scope requirements are "compliance symmetric". That is, they don't privilege one way of complying with them over any other—any way of satisfying the complex condition specified by a wide scope requirement is as good as any other, as far as the requirement itself is concerned. It's precisely this feature of wide scope requirements that enables them to avoid many of the objections to their narrow scope counterparts, since the latter (unlike the former) require a particular response. Narrow scope requirements are thus are vulnerable to counterexamples in which the response they require doesn't in fact seem to be required (the "too strong problem") and in which they require conflicting responses (the "conflict problem").

But it's a curious fact that the very feature that allows wide scope requirements to avoid many of the traditional counterexamples to narrow scope requirements—namely, their symmetry—is *also* what many have found to be problematic. And that's because not all ways of complying with wide scope requirements are intuitively on a par. Recall (ME-Wide):

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

As various authors have pointed out, while intending means *M* because you believe *M* is a necessary means to your end *E* would be a perfectly rational response, dropping your belief that *M* is a necessary means to *E* because you intend *E* and don't intend *M* seems terrible, and not in any way rational. Nor does it seem rational for you not intend *E* just because you don't intend *M*.[159] We might expect this asymmetry to be captured by the requirement governing means-end coherence, and yet (ME-Wide) fails

[159] For different versions of this complaint, see Kolodny (2005), Schroeder (2009), Way (2010), and Lord (2014).

miserably at such a task—unlike (ME-Narrow), there is more than one way to comply with (ME-Wide), and it doesn't privilege any one way of complying over any another.

The same compliant arises for (AA-Wide):

(**AA-Wide**) Rationality requires of you that you not: believe you ought to φ and not intend to φ.

Whereas intending to φ because you believe you ought to seems like a perfectly rational response, giving up your belief that you ought to φ because you don't intent to φ seems terrible. Yet (AA-Wide) is silent about why—as far as its concerned, either way of not being akratic is equally good. And so it turns out that the wide scope view's greatest strength and main attraction—its symmetry—is also its main weakness and source of resistance.[160]

I mentioned in Chapter 2 that wide scopers like Broome are sensitive to this complaint, but don't view it as an objection to (ME-Wide) or (AA-Wide) themselves since they don't think it's the job of the principles so expressed to explain everything that might go wrong in such cases. Instead, they take the lesson to be that such principles need supplementation. Broome's preferred route is to introduce a special class of diachronic "basing prohibitions" and "basing permissions" that rule out problematic transitions (or basing relations more generally) between attitudes and permit acceptable ones. Although doing so comes at the cost of complexifying the wide scope view, Broome (2013) rightly points out that

[160] The emphasis on the impermissibility of various "transitions" like those above is a major reason why most narrow-scopers—including Kolodny, Lord, and Schroeder—construe narrow scope requirements diachronically. But narrow-scopers who prefer synchronic requirements can raise a similar complaint.

narrow-scopers *also* need to appeal to such principles to explain the full range of our intuitive verdicts concerning when and whether a given transition between mental states is rationally permissible.

What's notable is that the pressure-based view can explain the apparent asymmetry and do so *without* positing special-purpose basing prohibitions and/or permissions in the structural domain, and without having to go "diachronic". The fact that the direction of attitudinal pressure reflects that of justificatory pressure, together with the fact that *not* having an attitude doesn't itself generate attitudinal pressure, can all the work needed. For suppose you intend end *E*, believe *M* is a necessary means to *E*, and yet do not intend *M*. On the pressure-based view, your intending *E* together with your belief that *M* is a necessary means to *E* generates substantial (pro tanto) attitudinal pressure to intend means *M*, just as your having good (pro tanto) reason to intend *E* together with your having good reason to believe that *M* is a necessary means to *E* gives you good reason—i.e. ensures that there is substantial (pro tanto) justificatory pressure—to intend *M*. In the absence of conflicting attitudes is would thus be structurally irrational for you not to intend *M*, just as we'd expect.

In contrast, your intending *E* together with your not intending *M* does *not* generate any attitudinal pressure to revise your belief that *M* is a necessary means to *E*, just as your merely having good reason to intend *E* does not give you good reason—i.e. does not ensure there is any justificatory pressure—to revise your belief that *M* is a necessary means to *E*. And so it would be structurally irrational for you to revise your belief, which is (again) the intuitively correct verdict. Similarly, your merely *not* intending *M* together your belief that *M* is a necessary means to *E* fails to generate any attitudinal pressure to give up your end *E*, just as your having good reason to believe that *M* is a necessary means to *E* is (by itself) simply silent on whether you have good reason pursue, or instead give

up, your end *E*. And so to give up your end *E* solely on the basis of your not intending *M* together with your belief that *M* is a necessary means to *E* would be irrational, as desired.

It's only by actually *adopting* some attitude towards means *M* that any attitudinal pressure to drop end *E* can be generated. Suppose, for example, that although you intend *E* and believe *M* is a necessary means to *E*, you believe that *M* would involve doing something undesirable, and so you are reluctant to pursue it. In such a case your being negatively disposed towards *M* together with your belief that *M* is a necessary means to *E would* generate attitudinal pressure to give up your end *E,* and this pressure would conflict with the existing attitudinal pressure to take means *M*. To resolve the tension you face a choice: give up your end *E* or take the means *M*. Which way you should go will depend on the strength of your commitment to *E* as compared to to your reluctance to take *M* (which may ultimately result in an intention to *not M*), as well as what other attitudes you may have towards the situation. As expected, this is precisely analogous to a case where you have good (albeit *pro tanto*) reason to pursue *E* but also have good reason to not take the means *M* you have good reason to believe to be required—the reason to pursue *E* conflicts with the reason to not take *M,* and the force of both is weakened (or "defeated" to some extent) and it's not clear how you should proceed. It'll depend on the strength of the respective reasons, as well as on other facts about what you have reason to do.

A similar story can be told to explain our intuitive verdicts concerning the different ways of complying with (AA-Wide) and other wide scope requirements. The pressure-based view thus enjoys a significant explanatory advantage over both the wide scope and narrow scope views, both of which have to resort to special-purpose basing principles to explain the data above.

## 5.5 Differences between dimensions

Although I've dwelt on the similarities between the two dimensions of rationality, there are important ways in which they diverge. To begin with, 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, although 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 satisfy the antecedent of, say, Deductive Coherence by believing $p$ and believing $p$ entails $q$, 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) or justification, 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.

Notice that this account of how one can "escape" from attitudinal pressures is much simpler and more straightforward than the account of how one can "escape" from a requirement (or at least make it so that it no longer applies) that I outlined in Chapter 2. The account offered there appealed to the law-like status of requirements and the intuitive distinction between jurisdiction and conditions of application. Nothing similarly fancy is required on the pressure-based view—one "escapes" simply by revising one's attitudes.

In this respect the view on offer is similar to that of Sam Shpall (2013, 2014), who offers an account of the nature of rational "commitments" (as opposed to reasons or requirements) according to which they're grounded in an agent's attitudes and hence similarly escapable. His guiding overall

thought is that to be rationally committed to having a given attitude A is to be such that "you must be irrational if you fail to have A, assuming no changes in your other attitudes" (2014, 148). Thus, like requirements and unlike reasons, commitments are inherently strict. But unlike requirements and like reasons, commitments play a "non-decisive, or pro tanto, weighing role" in determining what one is required to do, all things considered (157). Despite both being pro tanto and escapable, however, commitments and attitudinal pressure are importantly different. Most obviously, whereas commitments are strict, attitudinal pressures are variable in force, depending on which other attitudes are present. What's more, whereas Shpall takes commitments to be normative in the sense that "they *put* genuine pressure on the committed agent to form the attitude to which he's committed" (149, my emphasis), on my view this gets things backwards: rational commitments are the result of the pressures generated by the agent's attitudes, not what generates the pressures in the first place. Rational commitments are at best derivatively normative. Most importantly, however, I think the very same sort of considerations that cast doubt on the existence of rational requirements in the strong sense can be marshaled against the existence of rational commitments in Shpall's normative-weight-bearing sense. Neither one is needed to do important explanatory work, and their plausibility can be explained in terms of attitudinal pressure.

A second important difference between the dimensions is that while it's plausible to think there will always be a fairly determinate verdict concerning whether—and, if so, to what degree—some 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 parallel verdict in cases involving mere attitudes. For in cases of attitudinal conflict, such as ones involving inconsistent beliefs, it'll typically be impossible for you to have a reasonable twin—that is, a person who has all the same attitudes but, unlike you, is justified in having them. And so the pressures generated by the conflicting attitudes won't combine as happily as

they do in cases where there is no attitudinal conflict, which are cases where the attitudinal forces combine in much the same way as the justificatory forces would combine were the attitudes to be justified. Thus, even if there's always a global "all things considered" verdict concerning what particular attitude(s) the combined justificatory pressures support, we shouldn't expect there to always be one concerning what the combined attitudinal pressures support. Instead, we'll sometimes have to be content with local "pro tanto" verdicts concerning the strength and direction of attitudinal pressure generated by particular (combinations of) attitudes, and the only "all things considered" verdicts to be had will be irreducibly disjunctive. A pair of examples will make the intended contrast clearer.

Suppose first that Tom is convinced he's Superman and that Superman can fly but hasn't previously 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 rational pressure to believe that he can fly and so, given the absence of conflicting attitudinal pressure, the most (structurally) rational response is for him to believe that he can fly. That's the response that's "called for" by his standing beliefs, and so 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 his flying runs contrary to all his evidence.

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 Superman's ability to fly. What then? Well, unlike the previous case, there doesn't seem to be a determinate 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 rational

pressure to revise the third—hence the sense in which, given any two of his beliefs, he should revise the third—and this inherent instability or conflict can only be resolved by revising at least one of the beliefs in question.[161] Thus, the only realistic way he can avoid flouting substantial attitudinal pressure is by escaping it, but no particular escape route is privileged. At most we can say that, all things considered, Tom is rationally required to *not* have all three beliefs together.[162]

Call the former case the "Coherent Tom" case and the latter the "Incoherent Tom" case. Taken together, such cases show that *even when evaluated on their own terms*, neither the wide-scope nor the narrow-scope view of structural rationality is fully satisfactory. In particular, the narrow-scope theory delivers intuitively incorrect verdicts in cases of attitudinal *conflict*—these are cases in which structural rationality requires something irreducibly disjunctive of you, just as the wide-scoper insists—while the wide-scope theory delivers intuitively incorrect verdicts in cases of attitudinal *concord*—these are cases in which structural rationality requires something specific of you, just as narrow-scopers insist. In other words, although attitudinal pressures sometimes combine in such a way as to require, or at least recommend, the avoidance having a bad combination of attitudes—such as the case of Incoherent Tom—other times they combine in such a way as to require or recommend the having (or avoidance) of particular attitudes—such as the case of Coherent Tom. The "directedness" of attitudinal pressure and the variability in its verdicts is something that's obscured by the focus on requirements, together with

---

[161] Granted, most actual situations are far more complex, and someone like Tom will typically have a large number of relevant attitudinal pressures. But even if the overall pressure usually militates in one direction, the point is that it doesn't (or at least needn't) always do so.

[162] This is of course consistent with the various justificatory pressures singling out a particular attitude for revision; but that's neither here nor there. As (foot)noted at the outset, I don't intend to take a stand on how we should think of agents who perform well along one dimension but less well along another. In particular, I won't be addressing the issue of whether it makes sense to compare performances along each dimension—other than maximal ones—and thereby form a judgment as to whether (or to what degree) such an agent is rational *tout court*.

the assumption that they're all similarly structured. Accordingly, even if we pretend like there's an interesting debate concerning whether the 'ought' or 'rationality requires' operator (under some suitable interpretation that doesn't render the debate empty) takes wide or narrow scope over, say, a material conditional operator, neither option does justice to all our intuitions concerning what structural rationality requires in various circumstances, and why.

## 5.6 The property of rationality

At this point it might help to say something more about the property of being fully rational. For as Broome (2007a: 363) notes, "[p]erhaps the most important question a system of rational requirements needs to settle is whether you are [fully] rational—have the property of rationality." For requirements-based theorists like Broome the answer is (deceptively) straightforward: you are fully rational if and only if you satisfy all the requirements of rationality that apply to you. Of course, since I reject the requirements-based account of rationality, I don't face the question Broome considers. Nonetheless, an analogous and equally important 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 reasons rationality as well. So Broome's question bifurcates—there are two kinds of rational pressure (justificatory and attitudinal) and two corresponding properties (being reasons rational and being structurally rational), and we need to say something about the relationship between each. Thus:

(Q1) What's the relationship between justificatory (or J-)pressure and the person-level property of being fully reasons (or R-)rational?

(Q2) What's the relationship between attitudinal (or A-)pressure and the person-level property of being fully structurally (or S-)rational?

It should be noted from the outset, however, that 'the property of being rational'—like 'the property of being moral'—harbors an additional ambiguity. In particular, it may refer to a property a person has in virtue of having a certain capacity, or instead a property a person has in virtue of attaining something. As Broome (2013: 110) notes,

> If you say of human beings in general that they are moral animals, you probably mean they have a capacity for morality. The opposite of 'moral' in this sense is 'nonmoral'. Similarly, to say that human beings are rational animals is to say they have a capacity of rationality. The opposite of 'rational' in this sense is 'nonrational'. When I use 'rational' and 'moral' [I] refer to attainments. The opposite of moral in my sense is 'immoral'; the opposite of 'rational' is 'irrational'.

Broome's question concerns the property of being fully rational, where being fully rational in the relevant sense is (i) a property of persons, (ii) all-or-nothing—a person is either fully rational, or they're not—and (iii) an attainment, not a capacity.[163] But we also talk about (sets of) individual attitudes being rational, and of some attitudes being more rational than others. This is something a theory of rationality must accommodate as well.[164]

So we should distinguish the *person*-level property of being rational and the *attitude*-level property of being rational, as well as the *all-or-nothing* property of being fully rational and the *graded* property of being rational to a certain (possibly non-maximal) degree. Since these distinctions cross-cut

---

[163] The capacity is obviously important; it's just not of present concern.

[164] Cf. Reisner (2009).

each other, we end up with four relevant properties (all in the sense of an attainment): an all-or-nothing person-level property, a graded person-level property, an all-or-nothing attitude-level property, and a graded attitude-level property. What's more, once we distinguish reasons rationality and structural rationality, things get even more complicated—the four properties cross-cut the two dimensions, bringing the total to eight. (And, as we'll see, there are others.)

It's an open question how the different properties are related. And for the purposes of this paper I needn't take a definite stand—the pressure-based account of rationality is compatible with a variety of answers. However, for the sake of concreteness I'll briefly sketched what I take to be the most initially promising proposal.

Let's begin with the all-or-nothing properties. At the outset of the dissertation I claimed that to be fully rational *simpliciter* is to have a fully justified, coherent set of attitudes. Being fully rational *simpliciter* is thus an all-or-nothing person-level property, but it's one that is understood in terms of two others: for a person to be fully rational *simpliciter* is for them to be fully reasons (or R-)rational and fully structurally (or S-)rational. (Henceforth I'll often drop the 'fully' qualifier; it should be understood as implicit until I turn to the graded properties.) The latter two all-or-nothing person-level properties are, in turn, understood in terms of all-or-nothing set-of-attitude-level properties: for a person to be R-rational is for them to have a justified set of attitudes and for them to be S-rational is for them to have a coherent set of attitudes. What is it for a *set* of attitudes to be justified or coherent—and hence for it to be (in my terminology) R- or S-rational? Arguably, it's just for each *individual* attitude in that set to be justified or coherent. A *person* will thus be R- or S-rational just in case each of their individual attitudes are R- or S-rational. On this view, to be rational (in either sense) is a matter of having individually rational attitudes—a person is rational just in case each of their relevant attitudes is rational.

In the case of structural rationality, however, this might seem wrong. Isn't structural rationality first and foremost a property of *groups* of attitudes, and only derivatively a property of the individual attitudes belonging to such groups? According to the requirements-based view, it is; according to the view I favor, it's not. For we need to distinguish carefully between the facts in virtue of which an *individual attitude* is S-rational and the facts in virtue of which a *set of attitudes* is S-rational. Everyone can agree—pressure-based and requirements-based theorists alike—that the facts in virtue of which an attitude is S-rational include facts about one's other attitudes. But that's *not* the same thing as saying that an individual attitude is S-rational in virtue of belonging to a set of attitudes that is S-rational, as the requirements-based view would have it.[165] So even though the individual attitude-level property of being S-rational obtains in virtue of facts about one's other attitudes—and in particular how well it coheres with (or, in my terminology, is "attitudinally supported" by) them—it doesn't follow that it obtains in virtue of the set-of-attitudes-level property of being S-rational obtaining. In general, we shouldn't confuse facts about the *X*s with facts about the set of *X*s. It's therefore possible—and, given the pressure-based view, plausible—to understood the set-of-attitudes-level property in terms of the individual-attitude-level property with respect to *both* reasons rationality *and* structural rationality. The two dimensions merely differ in terms of what grounds or explains their respective individual-attitude-

---

[165] Reisner (2009), like Broome (2007, 2013) is explicit about the primacy of the sets-of-attitudes-level property, though he recognizes there are alternatives:

> "I have made the [primary] object of rational evaluation [in the case of structural rationality] the set of an agent's mental states, or some proper subset thereof, but there are alternatives: a vexing question is whether the proper objects of rational evaluation are collections of mental states or of the agents who hold them."

Reisner doesn't consider the possibility of taking the primary object of rational evaluation to be an agent's *individual* mental states, but that's to be expected—he's operating within the standard requirements-based framework, which is invariably concerned *combinations* of attitudes.

level properties—it's our other attitudes in the case of structural rationality and our reasons (or evidence or whatever) in the case of reasons rationality.

If this is right, then the relationship between rational pressure and the all-or-nothing person-level property of being rational boils down to the relationship between rational pressure and the all-or-nothing attitude-level property of being rational. Our original questions thus become:

(Q1*) What's the relationship between J-pressure and the all-or-nothing attitude-level property of being R-rational?

(Q2*) What's the relationship between A-pressure and the all-or-nothing attitude-level property of being S-rational?

The answer, I'll suggest, is plausibly the same in both cases.

Let's begin with reasons rationality, since there's already a substantial literature devoted to the question (under various guises) of the relationship between J-pressure (reason, justification, warrant, etc.) and an attitude having the property of being R-rational (reasonable, justified, warranted, etc.). Here it's standard to begin by drawing yet another important distinction—one between two all-or-nothing attitude-level properties. In epistemology, for instance, it's common to distinguish between a belief's being "propositionally" or "ex ante" justified and it's being "doxastically" or "ex post" justified.[166] This corresponds roughly to the distinction between having *justification to believe p* and having a *justified belief in p*. As Pryor (2001: 104) notes:

[166] The ex ante/ex post terminology, which I'll adopt, is due to Goldman (1979).

> [S]ometimes when we're evaluating your epistemic standing, we're only interested in whether you have justification for believing certain *propositions*—regardless of whether you actually do believe those propositions. Other times, though, were interested in more than that. We want to know whether you *do believe* the propositions you have justification for believe, and if so, whether your belief is *based on* that justification. You may have very good reasons for believing *p*, but base your belief in *p* on bad reasons. In such cases, your *belief* is epistemically defective, even though it's a belief in a proposition you have justification for believing.

It's thus generally agreed that being ex post justified in believing *p* requires not just that one be ex ante justified in believing *p*, but also that one's belief in *p* be "well-founded" or "properly based", where this amounts to being formed (and/or sustained) in an appropriate way in response to that which provides justification to believe *p*—namely, one's reasons or evidence.[167,168] It's thus possible to be ex ante justified in believing *p* without being ex post justified—one might believe *what* the evidence supports without doing so *because* the evidence supports it, and instead simply as a result of (say) a cognitive bias.

We thus need to distinguish between two all-or-nothing attitude-level properties: that of being ex ante justified and that of being ex post justified. This distinction applies to the practical side of reasons rationality just as much as it does the epistemic side.[169] A given action or non-doxastic attitude (intention, desire, preference, etc.) can be ex ante justified without being ex post justified, just as a doxastic attitude can. As it's often put, you might (intend to) do something you have good reasons to do, but those reasons might not be the reasons *for which* you (intend to) do it. Being ex post justified in

---

[167] See Feldman and Conee (1985) for more on the contrast between "well-founded" and "ill-founded" beliefs. For an argument against the need for proper basing, see Silva (forthcoming-a).

[168] Objections have been raised to the orthodox view of the relationship between *ex ante* and *ex post* justification, including by Turri (2010). See Silva (forthcoming-b) for a response.

[169] And of course it applies to morality as well. In T.S. Elliot's famous words: "The last temptation is the greatest treason, to do the right thing for the wrong reason".

$\varphi$-ing requires not just that one be ex ante justified in $\varphi$-ing, but also that one's action or attitude $\varphi$ be properly based.

The ex ante/ex post distinction arises with structural rationality as well. Just as one can have an attitude that is ex ante R-rational (or justified) without being ex post R-rational, so one can have an attitude that is ex ante S-rational (or coherent) without being ex post S-rational. Suppose Adam is a die-hard conspiracy theorist with a vast web of coherent (though largely unjustified) conspiracy-related beliefs. One of the things Adam believes is that 9/11 was an inside job. This belief is attitudinally supported by—and hence "coheres" with—a large number of other beliefs he has, including that George W. Bush and Dick Cheney desperately wanted an excuse to invade Iraq, topple the regime, and take control of Iraq's oil supply. But suppose that Adam's belief that 9/11 was an inside job isn't in fact based on any of the beliefs that provide the relevant attitudinal support—it's not a belief he has reasoned to or inferred from his other conspiracy-related beliefs. Instead, it's simply a byproduct of his deep-seated hatred of the United States and its elected officials. Because of this hatred, Adam always assumes the worst of American politicians, and then—if needed—reasons (or rather "reasons") his way backwards to arrive at a coherent set of beliefs. Adam excels at ad hoc rationalization. So although he ends up with a set of beliefs that attitudinally support the belief that 9/11 was an inside job, he doesn't believe the latter *on the basis* of the former. It's a belief that is ex ante S-rational, but not ex post—it's a coherent attitude for him to have, but it's not based on what *makes* it coherent.

With respect to both dimensions of rationality, then, it's plausible that in order to be ex post rational an attitude has to be not just ex ante rational, but also properly based. What it is for an attitude (or action) to be properly based, and what exactly it can and should be based *on* in either case, are vexed

questions.[170] Fortunately we needn't delve into the details. And that's because the pressure-based account of structural rationality, like that of reasons rationality, is compatible with a variety of different accounts of the basing relation. All that matters for present purposes is that with respect to *both* dimensions of rationality we can—and should—draw a distinction between being ex ante rational and being ex post rational, both of which are all-or-nothing attitude-level properties, and that in each case the ex post property is plausibly understood (partly) in terms of the ex ante property.

Taking stock: to answer Broome's original question we needed to say something about the relationship between rational pressure and the all-or-nothing *person*-level property of being rational. I then suggested this question boils down to the relationship between rational pressure and the all-or-nothing *attitude*-level property of being rational. The ex ante/ex post distinction, however, reveals the need for further refinement. Since to be ex post rational an attitude has to be ex ante rational (as well as properly based), the relevant question now becomes: what, exactly, is the relationship between rational pressure and the all-or-nothing attitude-level property of being ex ante rational? Once again, this question arises with respect to both dimensions of rationality. Thus:

(Q1**) What's the relationship between J-pressure and the all-or-nothing attitude-level property of being ex ante R-rational?

(Q2**) What's the relationship between A-pressure and the all-or-nothing attitude-level property of being ex ante S-rational?

---

[170] Korcz (2010) provides a useful overview of the "epistemic" basing relation. Analogous issues arise with the "practical" basing relation. The correct account of proper basing may very well differ between dimensions, though presumably there will be important similarities as well.

And, as before, I think the answer is plausibly the same in both cases.[171]

Let's again start with reasons rationality: how are we to understand the relationship between J-pressure (reason, justification, warrant) and an attitude's being fully ex ante R-rational (reasonable, justified, warranted, etc.)? Although there are a variety of possible answers, here's a standard one: for a given attitude to be ex ante R-rational is for it to be the attitude that is best supported by the totality of one's reasons—it's the attitude there is most J-pressure to adopt, all things considered.[172] In epistemology, it's a commonplace that insofar as what one is justified in believing depends upon one's evidence, what's relevant is the bearing of one's *total evidence*.[173] Following Carnap (1947) and Hempel (1960), this is known as "the requirement of total evidence", and something similar is standardly assumed when it comes to practical (as opposed to theoretical) reasons rationality—what's relevant is the bearing of all one's practical reasons, taken together. Similarly, I suggest, what's relevant for both practical and theoretical structural rationality is the bearing of all one's attitudes, taken together.

In the same way, I think it's plausible that for a given attitude to be ex ante S-rational is for it to be the attitude that is best supported by the totality of one's other attitudes—it's the attitude there is most A-pressure to adopt, all things considered. In neither case is an attitude's being ex ante 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 rational pressure to φ considered on its own—that is, to a proper

---

[171] It's worth emphasizing, though, that the pressure-based approach to rationality is compatible with a variety of different answers. I'm merely indicating what I take to be the most plausible view.

[172] The relevant kind of reason is *epistemic* in the case of epistemic reasons rationality and *practical* in the case of practical reasons rationality.

[173] Cf. Kelly (2008, 2014)

subset of one's reasons or evidence concerning φ in the case of reasons rationality or a proper subset of one's attitudes relevant to φ in the case of structural rationality.

The force metaphor familiar from Newtonian mechanics gives us another way of saying the same thing. For just as it's standardly thought that being reasons rational in φ-ing is a matter of responding correctly to the *resultant* justificatory pressure concerning φ generated by the totality of one's reasons or evidence, rather than to an particular *component* justificatory pressures generated by proper subsets of one's reasons or evidence, so it's plausible that being structurally rational in φ-ing is a matter of responding correctly to the resultant attitudinal pressure concerning φ generated by the totality of one's (relevant) attitudes, rather than to any particular component attitudinal pressures generated by proper subsets of one's attitudes.

I readily admit, however, that there are important and difficult questions concerning how the component rational forces combine in each case. Following Pryor (2004), I provided a heuristic in the case of structural rationality according to which 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, holding everything else fixed. (I called this the "Proportionality Thesis"). But this of course does little more than pass the buck to reasons rationality. I don't consider that an objection, however, since I neither need nor want to take a stand (at least for the purposes of the dissertation) on the exact shape of the "vectorial" function for either reasons rationality or structural rationality. Though important in its own right, it's an issue that is downstream from the one I'm presently concerned with, which is whether we should think of structural rationality as being fundamentally requirements-based (like the law) or instead fundamentally pressure-based (like reasons

rationality). A detailed account of the normative combinatorics may not be possible, any more than it is in, say, morality, but either way it's a topic for another day (or two or…).

So far I've focused on "threshold-y" (and hence all-or-nothing) properties, but of course we also talk about properties that come in degrees—we talk about some actions being *more rational* than others (as well as some people being more rational than others). This makes perfect sense on a pressure-based conception of rationality—a given attitude can be better supported than another, whether by one's reasons or instead by one's other attitudes, and hence more rational. Coherence, on this view, is fundamentally a gradable property that a given attitude can enjoy more or less of—it's a matter of how well it's positively supported by one's other attitudes. Similar explanations are available for more nuanced judgments involving comparative judgments, including that some attitudes are *substantially* (or slightly, etc.) more rational than others. Requirements-based accounts can also make sense of such talk, at least to some degree. For example, one attitude might be considered "more rational" than another just in case having it, holding all other attitudes fixed, would violate fewer requirements. And when we talk about some attitudes being (e.g.) substantially more rational than others, perhaps than can be understood as claiming that the former attitudes violate substantially fewer requirements than the latter ones. This is considerably less natural, however, and the pressure-based view enjoys the advantage of not having to rely on 'counting' (though it does have to rely on measuring or comparison).

To sum up: on the view I favor, the relationship between rational pressure and the all-or-nothing attitude-level property of being ex ante rational is the same with respect to both dimensions. And it provides the basis for a (partial) answer to Broome's original question concerning the all-or-nothing person-level property: on the pressure-based view, a *person* is rational only if they do (believe, intend, etc.) what there is most rational pressure to do, and do so in the right sort of way. Just as being reasons

rational is a matter of doing (believing, intending, etc.) what the 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 both cases the all-or-nothing person-level property is understood partly in terms of the all-or-nothing ex ante attitude-level property, which is in turn understood in terms of all-things-considered facts about rational pressure. And these global, all-things-considered facts about rational pressure are determined by the complex interactions between various local, graded facts about rational pressure and their grounds. It's these latter facts, then, that are normatively fundamental with respect to both domains of rationality.[174]

## 5.7 Underlying unity and the primacy of pressure

In closing, let me address an important worry otherwise sympathetic readers might have concerning the purportedly sharp distinction between kinds of rational pressure. Given the systematic parallels between justificatory pressure and attitudinal pressure, one might reasonably doubt that they are entirely distinct things, and instead hope to understand one in terms of the other, or both in terms of some third thing. Since what I have in effect claimed is that attitudinal pressure behaves in all the same ways that justificatory pressure does (aside some relatively minor differences), one might think it would be surprising if they were totally autonomous, and the obvious way to explain the similarities would be to try to show how they are ultimately unified.[175]

---

174 In both cases they must also do so in the "right sort of way"—what the ex ante/ex post distinction shows is that doing what there is most rational pressure to do isn't sufficient for being (fully) rational; at most it's necessary. (And I say 'at most' because there's a debate as to whether the correct model of rationality—and in particular practical rationality—is "maximizing" or instead merely "satisficing". I'm ignoring this debate; it would only serve to further complicate matters.)

175 Thanks to Jonathan Way for pressing this point.

As it turns out, however, this apparent weakness is in fact a strength, and arguably one of the main selling points of the view on offer. For although I take the dimensions to be distinct and autonomous, there is a way in which reasons rationality can be seen as enjoying a certain kind of explanatory priority over structural rationality. This is suggested by the observation, noted earlier, that facts about the amount (and direction) of attitudinal pressure generated by particular (combinations of) attitudes seem to track facts about the amount of justificatory pressure there would be were the subject to be justified in having those attitudes, holding everything else fixed. Structural rationality thus seems to behave as something like the "shadow" of reasons 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, and we're not always in a position to know which ones they are—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. So there does seem to be a sense in which reasons rationality enjoys explanatory priority, and in which there's an underlying unity between the two dimensions. What I remain skeptical of, however, is any attempt to reduce one to the other.

To sum up: on the view I've argued for, *pressure is primary*—that is, with respect to *both* dimensions of rationality, facts about rational pressure are normatively prior to facts about what one rationally ought, or is required to do. In both cases the relevant "threshold-y" ought-facts obtain in virtue of the corresponding "contributory" pressure-facts, and not vice versa. To say this is make an explanatory claim, of a broadly metaphysical sort. And so it's also correct to say that there's an important sense in which facts about rational pressure are *explanatorily* prior to facts about what one rationally ought to do. The former facts, taken together, metaphysically determine—and thereby explain

—the latter facts. This of course stands in marked contrast to the dominant view, which takes rational requirements (and principles more broadly) to occupy center stage in the theory of structural rationality, as well as to the tendency to take facts about what we rationally ought, or are required, to do as explanatorily basic.

Taking pressure to be primary opens up the possibility of rejecting the existence of any genuine ground-floor rational requirements, conditional or otherwise, other than the general constitutive requirements to be reasonable and to be coherent (or at least not be incoherent). That's all rationality 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 find ourselves in. I've also argued that it's a mistake to think that the difference between dimensions can be attributed, in whole or part, to the requirements-based nature of the structural rationality. The difference is instead traceable to the distinctive kind of rational pressure (attitudinal vs. justificatory) associated with each—pressures which in both cases should be understood as defeasible and holistic (or, better, regional), and which can interact in a variety of ways. Requirements are at best a footnote in the theory of rationality.

# CHAPTER 6: REASONS AND REASON

## 6.1 Reasons and *reason(s)*

The notion of a normative reason has played an increasingly prominent role in recent theorizing. This focus on reasons has been especially pronounced in ethics and metaethics, but more recently has taken hold in epistemology as well. The popularity of reasons-responsiveness as an account of rationality, both practical and theoretical, is just one manifestation of a much larger theoretical shift—one which takes the notion of a reason to be of central normative and theoretical importance.

Although I'm critical of this recent trend, I'm not going to be arguing directly against it in what follows. Instead, my goal will be to point out various subtleties in how we ordinarily think and talk about reasons—subtleties which, if taken seriously, have various upshots, both substantive and methodological. I'll focus on two subtleties in particular. The first concerns the use of 'reason' (in its normative sense) as both a count noun and a mass noun, and the second concerns the context-sensitivity of ordinary reasons-claims. The more carefully we look at the language of reasons, I'll argue, the clearer its limitations and liabilities become. The cumulative upshot is that although talk of reasons is intelligible and useful for the purposes of communication, we should be wary of placing much weight on it when engaging in substantive normative inquiry. By way of illustration, I'll consider some potential pitfalls of taking our talk of reasons too seriously, explaining how careful attention to the language of reasons undermines the main argument for moral particularism, Mark Schroeder's recent defense of Humeanism about practical reasons, and the "reasons-first" program in metanormativity.

I being this chapter by considering the various ways in which we standardly use the common noun 'reason', focusing on the relationship between the use of 'reason' (in its normative sense) as a

count noun and as a mass noun. I argue that there's a strong case to be made in favor of understanding reasons (count) in terms of reason (mass), rather than vice versa. I then explore some notable consequences of understanding reasons in terms of reason, and briefly assess its overall philosophical significance. In the following chapter I present new data concerning the contextual variability of reasons-claims that reveals a tension between the theoretical role that reasons are supposed to play and the more practical role that they—or at least our claims about them—in fact play. I conclude in by exploring some consequences for particular issues in (meta)ethics. It turns out that claims about reasons are only ever the beginning, and never the end, of normative inquiry.

## 6.2 The language of reasons

Talk of reasons can be confusing—at least to the theorist—despite being utterly commonplace. That's because there are several distinct strains in such talk, and as we've already seen with Broome's "quick objection" in Chapter 1, failing to be sensitive to their differences can easily lead one astray. To start with, like many others I think it's worth distinguishing between the reasons *why* something is the case (these are commonly called "explanatory reasons"), the reasons why—or *for which*—someone does something (commonly called "motivating reasons"[176]), and the reasons for someone *to* do something (commonly called "normative reasons").[177] Although this three-way distinction is widely recognized, there is considerable disagreement over the proper characterization of each class as well as the relationship between them. I'll consider the prospects of providing a unified analysis below. For now, however, I'll briefly mention two complications, before focusing on a third.

---

[176] Despite being standard, this label is misleading. We can presumably have "motivating" reasons for belief, for instance, but it's doubtful that we're actually motivated to have beliefs—at least in general. The same goes for other "reasons-responsive" attitudes (intentions, etc.).

[177] I'm ignoring the use of 'reason' to mean 'faculty of reason' and the use of 'reason' as a verb.

The first complication concerns the class of motivating reasons—those which help rationalize or explain (at the so-called "personal" level) why some agent *S* *φ*-s or is disposed to *φ*, where *φ* is a verb phrase denoting some action or attitude. For within the class of motivating reasons there's a further distinction to be drawn between (what I'll call) "factual reasons" and "teleological reasons", on the one hand, and what Davidson (1963) calls "primary reasons", on the other. The basic form of factual reasons-claims is 'S's reason for φ-ing is (was) that *p*' whereas the basic form of teleological reasons-claims is 'S's reason for φ-ing is (was) to ψ'. In both cases, however, there are other ways of saying more or less the same thing. What's crucial is that in both cases the reason specified by the relevant clause ('that'-clause in the former and 'to'-infinitival in the latter) is the content of one of the agent's action-guiding attitudes (e.g. belief or knowledge in the former and intention or action-guiding preference in the latter). Roughly put, a factual reason is a fact the awareness of which helps explain why the agent φ-s, whereas a teleological reason is a goal or purpose that an agent aims to achieve or promote by φ-ing and which thereby helps explain her doing so. Hence we say things like 'Victor's reason for leaving early was that he was tired' or 'Jen's reason for going into law was that it pays well' as well as 'Sasha's reason for becoming vegetarian was to reduce animal suffering' or 'The only reason (why) Karl exercised regularly was to lose weight'.

A primary reason, on the other hand, is a psychological state (e.g. belief, desire), or combination of such states, that helps explain—in the right sort of way—why an agent *φ*-s. Hence we say things like 'Jorge realized he was tired, and that's the reason (why) he left early' as well as 'Sally wanted to reduce animal suffering, and that's the reason (why) she became vegetarian' or 'Bob thinks he's Superman, and that's the reason (why) he thinks he can fly'. As the Victor/Jorge and Sasha/Sally examples make manifest, there's a close connection between factual/teleological reasons and primary reasons, and they

are plausibly seen as providing complimentary explanations rather than competing ones. Primary reasons-claims highlight the agent's contentful attitude(s)—those that play a certain explanatory role—whereas factual/teleological reasons-claims highlight the content of those attitudes.

The second complication concerns the class of normative reasons—the considerations which "count in favor of" performing certain actions and/or having certain attitudes. For as noted in Chapter 1, many have insisted on the need to distinguish between the reasons "there are" for some agent to $\varphi$ and the reasons that agent "has" to $\varphi$, where only the latter are of direct relevance to rationality and its ilk.[178] (A similar distinction arises with talk of evidence.) Suppose, for example, that you have skin cancer but it has yet to be detected. Although the fact that you have skin cancer is a reason for you to visit the doctor, it's not a reason you intuitively *have*—it's not something you could be expected to take into account in deliberation, nor anything else that plausibly makes a difference concerning what it's reasonable for you to do. So although your failing to go to the doctor's may be unfortunate, it wouldn't be *irrational* so long as you remain ignorant of your condition.

I myself am happy granting something like the distinction between reasons there are and reasons had. But there are a lot of complications concerning how such a distinction should be understood, and how exactly (if at all) it manifests itself in ordinary thought and talk. Since these additional complications won't matter in what follows, I'll simply help myself to the intuitive distinction without negotiating the details. Although for convenience I'll mostly focus on claims concerning reasons there are rather than reasons had, both should be kept in mind.

Unlike the first two complications, the one I'm most concerned with has been largely (though not entirely) overlooked. It's the distinction between the use of 'reason'—in its normative, but not

[178] The former are often called "objective reasons" and the latter "subjective reasons"—see, e.g. Parfit (2011) and Schroeder (2007)—though I think those labels do more harm than good.

motivating or explanatory, sense—as a count noun ('Julie has many reasons to lie') and as a mass noun ('Julie has lots of reason to lie'), as well as between the facts that such nouns are used to report.[179] Intuitively, count nouns denote (classes of) "things" that are countable, and hence can occur with cardinal expressions ('one', 'two', 'three', …) and take plural form (*-s*), while mass nouns denote "stuff" that's not countable, and hence do not occur with cardinal numerals and are generally singular or unmarked for number.[180] However, the term 'stuff' can be misleading, and is far less apt for so-called "abstract" mass nouns ('information', 'freedom', 'reason', 'advice', …) than it is for more "concrete", substance-denoting mass nouns ('water', 'cheese', 'sand', 'beef', …). The same is true of 'thing', since count nouns vary considerably in the clarity or precision with which their referents or denotations are individuated, with abstract count nouns oftentimes being particularly poorly individuated—contrast 'cat', 'chair', and 'microscope' with 'cloud', 'detail(s)', and 'explanation'. This reinforces the point that the mass/count noun distinction is a grammatical one, having to do with morphosyntactic and semantic properties, and contrary to what the intuitive gloss above might suggest it doesn't (by itself) have metaphysical implications.[181] For instance, although 'bean(s)' is count and 'rice' is mass, they both

---

[179] As Payne and Huddleston (2002) note, the mass/count distinction best understood as an instance of polysemy (more than one semantically related sense for a single word), rather than homonymy (distinct lexical items that happen to be pronounced and spelled alike).

[180] Mass nouns only combine with cardinal numerals (and other broadly numerical expressions) indirectly, requiring the presence of a so-called "classifier expression" like 'piece(s) of' or 'gallon(s) of', as in 'one piece of cheese' or 'two gallons of water'.

[181] This point is further reinforced by the fact that there are nouns that straddle the mass/count divide, including so-called "fake" or "collective" mass nouns ('furniture', 'silverware', 'luggage', …) and "plural" mass nouns ('earnings', 'belongings', 'dues', …). Morphosyntactically, fake mass nouns *are* mass nouns, but they are interpreted in much the same way as count nouns (see Barner and Snedeker (2005) for experimental evidenc, and McCawley (1975), Rothstein (2010), and Pelletier (2012), among others, for more general discussion). Plural mass nouns take plural form (often obligatorily), but otherwise behave like mass nouns (for discussion, see Gillon (1992), Ojeda (2005), and Schwarzschild (2009), among others).

denote granular substances, and although 'knowledge' is mass and 'belief' is count, they both denote mental states. In general, we shouldn't draw metaphysical conclusions solely from linguistic data.

Note that some determiners combine with both count and mass nouns, including 'any':[182]

(1a) I don't have any water. (mass)

(1b) I don't have any reason to lie.

(2a) I don't have any dogs. (count)

(2b) I don't have any reasons to lie.

Similarly promiscuous determiners include 'some', 'the', 'no', 'this', 'that', 'what', 'more', and genitives such as 'my'. Other determiners, however, are wholly or predominantly restricted to one class of nouns, and hence more discriminating. For example, 'each', 'every', 'several', 'few', 'a/an', 'many', 'either', 'neither', and 'one' all select for count nouns, while 'much', 'little', 'enough', and the lack of a determiner typically select for mass nouns. Hence the difference between:

(3a) There is water, but not much/#many. (mass)

(3b) There is reason to lie, but not much/#many.

(4a) There are dogs, but not many/#much. (count)

(4b) There are reasons to lie, but not many/#much.[183]

---

[182] I use the term 'determiner' broadly to include quantifying expressions as well as definite and indefinite determiners.

[183] There are actually two readings of (4b) with 'much', one of which means 'There are reasons to lie, but not much reasons' and the other of which means 'There are reasons to lie, but not to lie much'. It's the former that's relevant. Similar remarks apply to (3b).

Although a lot more could be said about the mass/count distinction, the important point is just that the common noun reason in its normative—but not explanatory or motivating—sense is standardly used both ways. The relationship between the two uses has nonetheless been neglected, with the vast majority of theorists focusing on the count noun. This is surprising because many common nouns in English pattern both ways, and when they do there is usually a story to tell about the relationship between the two uses—one which makes clear which is to be understood in terms of the other. And it's important because once we notice that a robust mass/count distinction arises with respect to 'reason', we're faced with the question of which, if either, use should be taken as basic, as well as which facts—count-y facts about reasons or mass-y facts about reason—are most normatively significant.

## 6.3 Reasons and reason: count and mass

There are a variety of ways in which mass/count noun pairs can be related.[184] Consider 'cheese', for instance. Although it is standardly used as a mass noun to denote quantities of cheese ('How much cheese do you want?'), it is also used as a count noun to denote kinds of cheese ('What cheeses do you like?'). Same goes for 'bread', 'coffee', 'virtue', and many others. Call this class of mass/count pairs **K**, for *kind*.

Next consider 'beer'. Like 'cheese' it is standardly used as a mass noun to denote quantities ('How much beer do you want?') and as a count noun to denote kinds ('How many beers are on tap?'). So it belongs to **K**. But as a count noun it can also be used to denote conventional units or individual servings of beer ('How many beers do you want?'). So 'beer' also belongs to what I'll call category **U**, for

[184] Cf. Payne and Huddleston (2002) and Gillon (2012).

*units*.[185] And of course there are further ways in which mass/count noun pairs can be related. Here are four of the main ones:

| | **Mass noun** | **Denotation of count noun** |
|---|---|---|
| **K** | *cheese, beer, virtue, activity* | **kinds** (*cheeses, beers, virtues, activities*) |
| | "I ate a lot of cheese." | "I've tried many different cheeses." |
| **U** | *beer, coffee, cake, pizza* | **units** (*beers, coffees, cakes, pizzas*) |
| | "I drank too much beer." | "I bought too many beers." |
| **I** | *detail, thought, action, error* | **instances**[186] (*details, thoughts, actions, errors*) |
| | "It contains much detail." | "It contains many details." |
| **S** | *sorrow, pleasure, light, anxiety* | **sources** (*sorrows, pleasures, lights, anxieties*) |
| | "It filled me with sorrow." | "Life is full of sorrows." |

So what about reason? There's a good case to be made that it patterns like the nouns belonging to S, all of which license (what I'll call) "generational" readings of verbs like 'give' and 'produce'. To see what I mean, consider the following pairs:

(5a) Taking warm baths is a simple pleasure of mine.

(5b) Taking warm baths gives me great pleasure.

(6a) The candle was the brightest light in the room.

[185] Notice that 'cheese' also belongs to **U**. The categories are not meant to be mutually exclusive—it's quite common for a common noun to belong to multiple categories.

[186] I realize this isn't happy terminology, since 'instance' is too broad—e.g. conventional units of *X* are also instances of *X*. So I'm using it quasi-stipulatively in a way that is suggestive rather than definitive of the relevant family. And it really is like a family, as opposed to a well-defined category. The difference between (e.g.) a **S-** and **I-** interpretation of a given expression is sometimes unclear. But just because there is overlap between families in some cases doesn't mean there is overlap between families in all cases.

(6b) The candle produced lots of light.

(7a) That we won fairly was a reason to celebrate.

(7b) That we won fairly gave us reason to celebrate.

Notice that in the first sentence of each pair above, the relevant count noun ('pleasure(s)', 'light(s)', 'reason(s)') is being used to describe some (type of) thing—whether it be an activity, event, fact, or something else—and in the second sentence that same thing is being described as explanatorily responsible for (or a "source" of) the stuff denoted by the mass noun ('pleasure', 'light', 'reason'). This pattern unifies the family of mass/count pairs belonging to **S**.

This isn't to say that members of **S** are *only* members of **S**—both 'sorrow' and 'light', for instance, belong to **S** but are also members of **K** and **I**, and hence can be used as count nouns to denote kinds as well as something like instances of sorrow and light, respectively. And of course there are many other members of **S**, including normative notions besides 'reason' (e.g. 'merit'). As far as I can tell, though, 'reason' (in its normative sense) only belongs to **S**. Also, notice that there is little metaphysical unity among the kinds of "stuff" picked out by the mass nouns belonging to **S**—it includes stuff that is physical ('light'), experiential ('pleasure'), emotional ('sorrow'), and normative ('reason'). Similar diversity arises in other families of mass/count pairs, such as **K** and **I**. And as with mass nouns in general, it's a further question what, if anything, composes each kind of stuff. For although in the case of physical stuff (e.g. water, light) it's reasonable to expect an answer to the question of what it's composed of (e.g. $H_20$, photons), it's far less clear we should expect an answer in cases of nonphysical or abstract "stuff" (e.g. sorrow, reason). In answering such questions we have no choice but to proceed on a case-by-case basis.

Importantly, however, the pattern that unites the family of mass/count pairs belonging to **S** involves an important asymmetry: with respect to each mass/count pair belonging to **S**, the count noun —i.e. what it applies to—is properly understood in terms of the mass noun—i.e. what is given rise to, or otherwise explained—rather than vice versa. To illustrate: imagine meeting someone blind from birth who asks you to explain what light is, as well as what lights are. There are various ways you might reply, but one thing you won't (or at least shouldn't) do is offer an explanation of light in terms of lights—to say, for example, that light is that which is emitted by lights. In contrast, you'd be hard-pressed to explain what lights are without making reference to light—that is, to avoid saying that lights are things which emits light.[187] Or imagine meeting an insensate who has never experienced pleasure. He asks you to explain what pleasure is, and what pleasures are. This is undoubtedly a difficult task, but one thing it makes little sense to do is offer an account of pleasure in terms of pleasures—although it might be *true* to say that pleasure is a state that is generated by pleasures, it wouldn't be *explanatory* in the relevant sense. In contrast, it makes perfect sense to characterize pleasures in terms of pleasure—to say, for example, that pleasures are things which give rise to pleasure.[188]

This provides the basis for a family resemblance argument in favor of understanding reasons in terms of reason, rather than vice versa. For just as sorrows, pleasures, and lights are naturally understood to be things (= facts, states, events, objects, whatever) which, in certain circumstances, are explanatorily responsible for there being sorrow, pleasure, and light, so we should understand reasons

---

[187] Of course, this wouldn't help much unless you gave the person a rudimentary explanation of light—but that's as it should be. Also, I'm not denying that when you're trying to provide an explanation of what it is to be a light it might sometimes be *adventitious* to first describe them as things that (e.g.) typically have a certain shape, give off warmth, etc, before characterizing them in terms of what they emit—namely, light. The order in which we explain things needn't always follow the order of explanation itself.

[188] Once again this passes the explanatory buck to the mass noun—as it should be.

(for *S*) to $\varphi$ to be things which, in certain circumstances, are explanatorily responsible for there being reason to $\varphi$. (Note: I'll oftentimes omit reference to the agent ('for *S*') in reason(s)-claims out of convenience.) And just as it would be a mistake to characterize sorrow, pleasure, and light in terms of sorrows, pleasures, and lights, so it would be a mistake to characterize reason in terms of reasons—to say, for instance, that reason to $\varphi$ is that which is generated by reasons to $\varphi$. That would be getting things backwards. Instead, reasons are to be understood in terms of reason, together with the (admittedly vexed) notion of explanation: reasons to $\varphi$ are things which are (at least partly) explanatory responsible for there being reason to $\varphi$.

The family resemblance argument isn't knock-down, of course. For as with any family there are going to be differences as well as similarities among members of **S**, and perhaps the relative priority of reasons and reason is such a difference. At the very least, though, the foregoing considerations shift the burden of proof onto those who wish to understand reason in terms of reasons, since they'll need to offer an alternative (and non-ad hoc) story about the relationship between our use of 'reason' and 'reason(s)' that vindicates the priority of the latter. Without such a story, however, we should be open to the possibility of following the argument where it leads. In the next couple of sections I'll consider some consequences of doing so.

## 6.4 A unified account of reasons-talk

One immediate benefit of understanding reasons in terms of reason is that it allows us to provide a pleasingly unified account of both normative and non-normative uses of the count noun 'reason(s)'—one according to which all reasons are explanatory reasons, and essentially so. To be a reason (of any kind) just is to be something that helps explain something else. Different "kinds" of reasons—

normative, motivating, explanatory, and so on—merely differ in terms of what's being explained (the explanandum), how it's being explained (the kind or level of explanation), or what's doing the explaining (the explanans). For instance, motivating reasons differ from other explanatory reasons primarily with respect to the level of explanation, together with the nature of the explanans—to explain why an agent φ-s in terms of (the contents of) the agent's beliefs, desires, intentions, and the like is to provide an explanation at the so-called "personal" level that is distinct from, but compatible with, explanations at the sub-personal level (e.g. in terms of lower-level cognitive processes) or non-personal level (e.g. in terms of neural or neurophysiological activity). Normative reasons, on the other hand, are distinguished primarily in terms of the explanandum. A normative reason is something that helps explain why there is (some amount of) reason to φ. Since what normative reasons help explain is the presence of something mass-y (i.e. reason) rather than something count-y, this proposal also naturally accounts for the absence of mass-y counterparts of non-normative uses of the count noun reason(s)—the nature of the explanandum differs.

Among those who seek to provide a unified account of our use of 'reason(s)', the view on offer is most similar to those of John Broome (2004, 2013), Stephen Finlay (2014), and Schroeter and Schroeter (2009). All three take normative reasons to be facts which (at least help) explain other normative facts; they just disagree over what the other normative facts are. Whereas Broome takes normative reasons to be facts which play a certain role in explaining why some agent *ought* to *φ*, Finlay takes reasons to be facts that explain why it would be *good* (in some way, to some degree) if some agent *φ*s. And Schroeter and Schroeter take reasons to be facts that "make a contribution toward" (and thereby helps explain) an actions "all told rightness". These are important alternatives, so it's worth seeing where they fall short—a task I'll return to as an appendix to this chapter [—a version of which I'll send along prior to the

defense. As with the appendix to Chapter 4, it'll be optional reading]. Although each alternative has its own weaknesses, one mistake they all share is prioritizing reasons instead of reason.[189]

## 6.5 The possibility of "groundless" normative facts

Another consequence of understanding reasons in terms of reason is that it allows for the possibility of "groundless" normative facts. For although there will always be (at least some) reason to φ whenever there is a reason to φ, nothing I've said guarantees that there is a reason to φ whenever there is reason to φ. The fact that there is reason to φ might instead be groundless; it might not hold in virtue of any particular fact identifiable as a reason. (Compare the possibility of there being light that is not produced by any lights.)

I mention this possibility not to endorse it, but merely to point out that it's an open question at this point whether the following biconditional fails in the left-to-right direction:

(**RR**) There is reason to $\varphi \leftrightarrow$ there is a reason to $\varphi$.

I should note that I've been following Schroeder (2007) and others in assuming that claims of the form *there is a reason to φ* involve bona fide existential quantification, and hence that the following holds:

(**RR+**) There is a reason to φ iff there is some *r* such that *r* is a reason to φ.

Enoch (2011) notes the importance of (RR+) for Schroeder's view, but raises a worry:

---

[189] Broome (2013) does briefly consider the possibility of taking the mass noun to be prior to the count noun, though he ends up dismissing it. For criticism of Broome's positive account of reasons, see Brunero (2013) and Dancy (2015). For criticism of Finlay, see my (forthcoming).

> What if some actions (or some other things) enjoy a basic, ultimate kind of positive normative status? [T]hey are actions that are to be performed and there is nothing in virtue of which they are to be performed. (443)

The present point is that *even if* (RR+) is true, the possibility Enoch considers isn't ruled out.

I should also note that some hear claims about what there is reason to do as being stronger than claims about what there is *a* reason to do, and so balk at the claim that there is reason to $\varphi$ whenever there is a reason to $\varphi$. One might therefore worry about the right-to-left direction of (RR) as well. But the difference in strength is arguably just pragmatic, and traceable to the threshold-y yet unspecific nature of bare (i.e. unquantified) reason-claims. To say there is reason to $\varphi$ is to say that $\varphi$-ing enjoys a non-zero amount of normative support, but (semantically speaking) it's to remain silent concerning exactly how much. So it would normally be pointless or misleading to make such claims unless the amount of support was significant. (Compare what's normally communicated versus strictly entailed by unquantified claims like 'There is water on the floor'.) Adding modifiers like 'at least some', 'little'/'not much', 'lots of', 'sufficient', 'decisive', and so on makes the strength of the intended claim explicit, and I for one have no trouble making sense of claims like 'There is reason to $\varphi$, but not very much'.

Indeed, the same sort of pragmatic story that Schroeder (2007) offers to account for the oddness of a claim like 'You have a reason to eat your bicycle' despite its literal truth can be used to explain the apparent infelicity of 'You have reason to eat your bicycle'. What Schroeder calls the "negative existential fallacy" is just an instance of a broader "fallacy" involving both mass-y and count-y quantificational claims—e.g. claims like 'There's no water left' are typically acceptable even when a couple drops remain. And this is itself an instance of a broader "fallacy" involving categorization in general—e.g. claims like 'I'm not angry' are often acceptable even if the speaker *is* angry, but not to a contextually salient degree.

## 6.6 Against the "reasons first" program

So far in this chapter I've played a bit fast and loose with various linguistic/conceptual/metaphysical distinctions, resulting in them being blurred. Such blurring often arises in philosophy—most obviously with questions concerning what it is to "count as" a thing of a certain kind. Talk of "counting as" sounds linguistic/conceptual but talk of being "a thing of a certain kind" sounds metaphysical. Questions concerning what it is to count as a thing of a certain kind are perhaps best understood as an exercise in something like what P.F. Strawson (1959) calls "descriptive metaphysics" or—better—what Emmon Bach (1986) calls "natural language metaphysics". The extent to which the broadly linguistic and conceptual argument I've given can be parlayed into a purely metaphysical (and not merely "descriptive" or "linguistic") one isn't clear, since it depends on thorny issues having to do with the relationship between language, thought, and reality—issues that I'm not able to address at any length here.

Nonetheless, at least in practice many philosophers who are fond of reasons proceed as if there's not much distance between our ordinary normative judgments involving the count noun 'reason(s)' and the underlying normative facts, taking it for granted that claims of the form '*r* is a reason for *S* to *φ*' are true just in case it's a fact that *r* is a reason for *S* to *φ*—or, in more theory-laden terms, just in case *r* has the relational property of being a reason for *S* to *φ*. Some are explicit about the presumed correspondence, but most are at least implicitly committed to it in virtue of their unabashed appeal to ordinary judgments about reasons when engaging in substantive normative (and metanormative) inquiry. If one takes our normative judgments involving the count noun 'reason(s)' seriously, however, it's hard to see why one wouldn't also take our judgments involving the mass noun 'reason' seriously—indeed, theoretical and methodological consistency would seem to require it. This is the dialectical space

within which the family resemblance argument in favor of understanding reasons in terms of reason gains the most theoretical traction.

Notice that if we do end up applying the considerations above concerning the way we deploy 'reason' and 'reason(s)' in ordinary thought and talk to our normative metaphysics, we'll be committed to rejecting both of the following:

**Reasons Primitivism**: The property (and concept) of being a normative reason is unanalyzable, and hence primitive.[190]

**Reasons Fundamentalism**: Facts about reasons—i.e. facts expressed by (true) claims of the form '*r* is a reason (for *S*) to *φ*'—are normatively fundamental.[191]

Note that facts *about* reasons are to be distinguished from the facts *which are* reasons—i.e. the facts corresponding to *r* in claims of the form '*r* is a reason (for *S*) to *φ*'. Whereas facts about reasons are normative facts, the facts which are reasons needn't be—and usually aren't.[192]

Although Reasons Primitivism and Reasons Fundamentalism are not often explicitly distinguished, it's worth keeping them apart. The former concerns analyzability, after all, while the latter concerns normative fundamentality, and neither entails the other. Schroeder (2007), for instance, accepts Reasons Fundamentalism but rejects Reasons Primitivism, offering a reductive analysis of facts

---

[190] See Scanlon (1998: 17), Skorupski (2010: 2), and Parfit (2011: 31) for claims to the effect that our concept of a normative reason is primitive. Presumably, however, they would say the same thing about the corresponding property.

[191] See, for instance, Raz (2000: 67), Schroeder (2007: 81), and Scanlon (2014: 2).

[192] Some have neglected this distinction; see Olson (2009) for critical discussion.

about reasons. Normatively fundamental facts needn't be fundamental *full stop*. Nonetheless, many philosophers who are attracted to one of the theses are also attracted to the other, and together they lie at the heart of the increasingly popular "reasons first" program in metaethics (or rather, metanormativity). Not everyone agrees on the exact contours of the program, of course, and there are various ways in which the core claims might be precisified. For example, we might distinguish "weak" and "strong" versions of Reasons Primitivism, where weak versions merely claim that the property (or concept) of being a normative reason is primitive while strong versions claim that it's the *only* normative property (or concept) that is primitive. Similarly, we might distinguish "weak" and "strong" versions of Reasons Fundamentalism, where weak versions claim that the fundamental normative facts *include* facts about which things are reasons to do what in which circumstances, while strong versions claim that the fundamental normative facts are *exhausted by* such facts. I should note that there is also disagreement concerning the arity of the reasons-relation. Scanlon (2014), for instance, takes it to be a four-place relation holding between a fact, agent, action, and set of circumstances, whereas Schroeder (2015) argues that at least one of its argument places is redundant. It's perhaps most common to treat it as a three-place relation between a fact, agent, and action, and for concreteness I'll be following suit. But, importantly, these sorts of in-house disputes won't matter so long as it's agreed that Reasons Primitivism and Reasons Fundamentalism concern the notion of a normative reason operative in ordinary thought and talk, since that's what the broadly linguistic considerations I've adduced (and will adduce) speak to.

It should be clear by now why those who take the linguistic data and family resemblance argument above seriously should reject Reasons Primitivism. For such theorists will be in a position to answer Scanlon's (2014) question of "whether there is something further to be said about what it is to

be a reason…something further that might be said to identify the property signified by that [relational] concept" (44). The answer will be 'yes': a given fact *r* has the property of being a reason (for *S*) to *φ* just in case—and *because*—*r* plays a certain role in explaining why there is reason to *φ*. The property of being a reason (and corresponding relation)—as well as the rest of properties corresponding to the count nouns in category **S** ("sources") above, such as the property of being a light, being a pleasure, and so on—will thus be analyzable in broadly functional terms, in much the same way that properties such as being a cook and being a gene are. To be a cook, it seems, is to play a certain (at least intended) role in the production of food, and to be a gene is to play a certain role in encoding and transmitting genetic information.[193] In general, a functional property is a property something has in virtue of the role it plays, whether that role be causal, teleological, computational, metaphysical, normative, or whatever. The details of the relevant roles are of course complicated, and functional analyses are rarely, if ever, expressible as a set of individually necessary and jointly sufficient conditions. But that's true of nearly all concepts and properties, and it's compatible with the claim that some concepts and properties can be analyzed partly in terms of others. If it wasn't, then nearly all concepts and properties would be on a par, and the corresponding forms of primitivism and fundamentalism would come cheap.[194]

It should also be clear why such a theorist should reject Reasons Fundamentalism. For if the property of being a reason is a functional property defined partly in normative terms, then facts expressed by true claims of the form '*r* is a reason (for *S*) to *φ*' will obtain partly in virtue of other

---

[193] Cf. Kim (1998: 25).

[194] Functional properties are hardly the only properties susceptible to analysis. To take just one example: so-called "conjunctive" properties, such as that of being a red barn, are commonly thought to be analyzable in terms of their constituent properties. To be a red barn, it seems, is just to be red and to be a barn (cf. Schroeder (2007)).

normative facts—in particular, (mass-y) facts about reason—and so not be normatively fundamental. Facts about reason will be normatively prior to facts about reasons, contra Reasons Fundamentalism.

## 6.7 An ecumenical middle ground?

By taking the mass/count distinction seriously, we've arrived at two preliminary conclusions: first, normative reasons (like all reasons) count as reasons in virtue of the explanatory role they play; second, what they help explain is some further (though closely related) normative fact. The first point suggests it's a mistake to take the property or concept of a normative reason as primitive, and the second suggests that it's a mistake to take facts about reasons to be normatively fundamental—they obtain partly in virtue of other normative facts.

Neither of these points may seem all that significant, however. And in many ways they're not. After all, to say there is reason to φ is just to say that there is a non-zero amount of normative support for φ-ing, and hence that φ-ing is to at least some degree worth doing or having. More colloquially, it's just to say that there is *something* (though not necessarily some *thing*) to be said in favor of φ-ing. This is little more than a paraphrase, though, not an analysis. And if that's all that can be said about what it is for there to be reason to φ, then it will turn out that Reasons Primitivists are at least partly right. Although they're wrong to take the property or concept of *a* normative reason to be primitive, they're right to insist that there's a positive pro tanto normative status in the vicinity that isn't amenable to further analysis. Replacing Reasons Primitivism with Reason Primitivism can be seen as simply shifting the focus to where it should have been (and to some extent was) all along.

A similar shift may appear plausible when it comes to Reasons Fundamentalism. Even if facts about reasons are not normatively fundamental, much of what Reasons Fundamentalists care about can

be preserved by switching to Reason Fundamentalism. Facts about reason, after all, are facts about a normative status that is pro tanto and non-strict—it's a status that can be "opposed" or "defeated" in various ways, and one doesn't automatically do anything wrong by failing to do what one has reason to do, since one may have more (or equal) reason to do something else. Unlike reasons, however, reason is a quantitative notion—it's something that comes in degrees or amounts, something you can have more or less of. So whenever someone makes a claim about there being reason to do something, there is always a further question concerning *how much* reason there is. This is what allows for comparative claims, whether between people ('Joe has more reason to lie than Sally, but I have the most') or between actions/attitudes ('Joe has more reason to lie than to hide'). The quantitative nature of reason also helps account for our talk of the "strength" or "weight" of reasons—the intuitive strength of a reason can be understood in terms of how much reason that fact is explanatorily responsible for, with possible amounts ranging from minute to decisive.[195]

Various distinctions between different "kinds" of normative reasons—practical, moral, epistemic, etc.—can be similarly preserved by understanding them in terms of the different kinds of reason—practical, moral, epistemic, etc.—they help explain. The same goes for the intuitive (but ultimately stipulative) distinction between the reasons there are to $\varphi$ and the reasons an agent has to $\varphi$—it can be understood as the distinction between the facts that help explain why there is a certain amount (and kind) of reason to $\varphi$ and the facts that help explain why an agent has a certain amount (and kind) of reason to $\varphi$.

It's in answering such questions—questions concerning kinds and quantities of reason—that our judgments about reasons are most relevant. After all, it's in virtue of helping explain the existence of

[195] It would therefore be a mistake to treat the difference between, say, a "decisive" reason and a merely "sufficient" reason as a difference in kind rather than degree.

reason that a fact counts as a (normative) reason at all. But we might go further and take reasons to play a *distinctive* role in explaining facts about reason: as suggested by the discussion in §6.2, we might take reasons to be what *generate* (and hence serve as "sources" of) reason. So even if facts *about* reasons are not normatively fundamental, the facts which *are* reasons might nonetheless play a privileged explanatory role—and one which would justify much of the attention that has been lavished on them in recent years.

Unfortunately, with respect to the last point, things aren't so simple. For as I'll argue in the next chapter, the facts we cite as reasons are rarely (if ever) themselves sources of reason. At best they're *partial* sources—on their own they're normatively irrelevant, and it's only in combination with other facts that any amount of normative support is generated. Again, that's at best. For it's at least arguable that some of the facts we unproblematically cite as reasons aren't even partial sources, but instead play a different explanatory role. (Examples of such will undoubtedly be contentious; I provide one in §7.3.)

More specifically, I'll present new data concerning the contextual variability of our judgments about reasons that reveals a tension between the theoretically substantive role that reasons-claims are supposed to play and the more pragmatic role that reasons-claims in fact play. For it is widely regarded as a truism that reasons are those things which "count in favor of" performing certain actions or having certain attitudes—they are "sources" of normative support (aka reason). What this amounts to is the assumption that reasons are things which *themselves* explain—in a broadly *metaphysical*, *non-pragmatic* sense—why there is reason to $\varphi$. But, as I'll argue, the things we actually cite as reasons only *partially* explain—in a broadly *communicative*, *pragmatic* sense—why there is reason to $\varphi$. As I'll put it, they're merely "representatives" of sources.

This contrast between roles rests on two cross-cutting distinctions: first, the distinction between full and partial explanations and, second, the distinction between pragmatic and non-pragmatic senses of 'explanation'. Both distinctions figure prominently in the literature on causal and scientific explanation. Since a detailed discussion of either is (unfortunately) beyond the scope of this dissertation, a few suggestive remarks will have to suffice.

The first distinction—between full and partial explanations—should be familiar enough. For almost all the explanations we actually provide are partial, not full. Suppose, for instance, that I order a triple espresso and you ask me why. I tell you it's because I'm exhausted. This is a perfectly sensible reply. Notice, however, that the fact that I'm exhausted does nothing, *by itself*, to explain why I order a triple espresso. It's only in combination with the fact that, say, I want to perk up, as well as the fact that a triple espresso will perk me up, and that I know (or at least believe) that a triple espresso will perk me up, that my reply comes close to fully answering your query. And even then there's significant room for further explanatory detail. The fact that I'm exhausted is thus only *part*—and in fact a rather small part—of the full explanation.

It's worth emphasizing that the provision of merely partial explanations is a general feature of our explanatory practice. It's true not just of action explanations, but also of all other explanations we provide as well—causal, scientific, historical, normative, metaphysical, you name it.

Talk of "providing" explanations brings us to the second distinction—the one between pragmatic and non-pragmatic senses of 'explanation'. Very roughly, we can think of pragmatic explanations as answers to "why"-questions, understood as the type of questions that form the core of our explanatory practice. Such questions are standardly posed in a particular context, by particular people—people who in asking the question already possess, and therefore presuppose, a huge amount of the background

information. A successful answer to such a why-question will be one that helps make sense of the phenomenon being asked about, or render it intelligible given the background information already possessed by the inquirer.[196] Slightly more fully, following James Woodward (2014) we can say that an explanation in the pragmatic sense is one that involves "irreducible reference to facts about the interests, beliefs or other features of the psychology of those providing or receiving the explanation and/or (ii) irreducible reference to the 'context' in which the explanation occurs". A *non*-pragmatic explanation is thus one that *doesn't* involve irreducible reference to such things. Instead, it solely concerns the objective explanatory relations—whether they be causal, nomic, metaphysical, or something else—obtaining between facts 'out there in the world'. A good pragmatic explanation will oftentimes be one that accurately tracks (some of) the objective explanatory relations involved, but due to the vagaries of context, the limitation of our knowledge, and various situation-specific demands, the relationship between the two will invariably be complicated.

Although a lot more could be said about both distinctions, all that's needed for present purposes is an intuitive grasp of each. The goal in the following chapter will be to gain a better understanding of the actual role of reasons-talk, thereby making the tension between it and its (at least purported) theoretical role more apparent.

[**Note**: As mentioned above, later this month I'll be sending along an appendix explaining and criticism the alternative accounts of reasons-talk provided by Broome (2004, 2013), Finlay (2014), and Schroeter and Schroeter (2009).]

---

[196] The notion of "making sense" is thus to be understood as implicitly relativized to particular inquirers or to the information they posses.

# CHAPTER 7: REASONS AND CONTEXT

In this chapter I present new data concerning the contextual variability of reasons-claims that reveals a tension between the theoretical (metaphysical) role that reasons are supposed to play and the more practical (communicative) role that they—or at least our claims about them—in fact play. I conclude in by exploring some consequences for particular issues in (meta)ethics. It turns out that claims about reasons are only ever the beginning, and never the end, of normative inquiry.

## 7.1 The variability of reasons

It has often been observed that what counts as a reason in one context may fail to count as a reason in another context, where the worldly (i.e. non-linguistic) facts are different. Reasons-claims thus exhibit *inter*-contextual variability. But what hasn't (yet) been recognized is that what counts as a reason in one context may fail to count as a reason *in that very same context*, where the worldly facts remain the same. Reasons-claims therefore also exhibit *intra*-contextual variability.[197] Notably, it's a form of contextual variability that also infects our thought and talk about causes, as Eric Swanson (2010) has argued. Indeed, much of my discussion in this section intentionally mirrors Swanson's discussion in order to make manifest the striking and systematic similarities between our thought and talk about reasons and our thought and talk about causes. I conclude the section by sketching an account of the role of reasons-talk that naturally accommodates the data.

---

[197] There are some technical notions of 'context' that individuate contexts extremely finely and so would count both types of variability as 'inter-contextual' variability. But it doesn't matter whether we decide for theoretical purposes to call the variability I'm concerned with 'intra-contextual' or not; what matters is just that it's different from the kind of variability usually appealed to, and more disconcerting for those who take reasons-claims seriously.

## 7.2 Some examples

Here's an example of the kind of intra-contextual variability I'm concerned with:

> **DANCE**[198]
>
> Billy loves to dance and there is going to be dancing at the party tonight. Billy, however, hasn't yet been informed of the party. I tell you about Billy and the party. In deliberating about whether Billy should attend, you might say any of
>
> (**D1**) The fact that there will be dancing at the party is a reason for Billy to attend.
>
> (**D2**) The fact that Billy loves to dance is a reason for him to attend.
>
> But it would be unacceptable to say any of
>
> (**D3**) #The fact that there will be dancing at the party is a reason for Billy to attend, and so is the fact that he loves to dance.
>
> (**D4**) #The fact that there will be dancing at the party and the fact that he loves to dance are both reasons for Billy to attend.
>
> (**D5**) #Since there will be dancing at the party and Billy loves to dance, there are two reasons for Billy to attend.

Claims like (D3)-(D5) fail to appropriately describe the stipulated normative relations. The dominant reaction among those I've consulted is that such claims are guilty of something like "double-counting". I think that's exactly right. Summarily put, what's wrong with (D3)-(D5) is that they double-count a source of normative support for Billy's going to the party. (I'll expand on this below.)

---

[198] This is inspired by one of Schroeder's (2007) favorite examples, though I'm using it for a very different purpose. For a similar example involving causation, see Swanson (2010).

It's worth noting that although DANCE concerns reasons *there are* (or "objective reasons"), the same variability arises with reasons one *has* (or "subjective"reasons"). For instance:

**SMOKE**

Sally opens the door to find the hallway filled with smoke. In explaining what grounds she has to believe there's a fire you might say either of the following:

(**S1**) That fact that she sees smoke is a reason for Sally to believe there's a fire.

(**S2**) The (known) fact that smoke is a sign of fire is a reason for Sally to believe there's a fire.

But it would be unacceptable to say any of

(**S3**) #The fact that she sees smoke is a reason for Sally to believe there's a fire, and so is the fact that smoke is a sign of fire.

(**S4**) #The fact that she sees smoke and the fact that smoke is a sign of fire are both reasons for Sally to believe there's a fire.

(**S5**) #Sally sees smoke and knows that smoke is a sign of fire, so she has two reasons to believe there's a fire.

As before, reasons-claims like (S3)-(S5) fail to appropriately describe the stipulated normative relations. They too are guilty of double-counting a source of normative support, though in this case it's a source of *epistemic* support rather than practical support. The variability I'm concerned with thus cross-cuts the objective/subjective divide as well as the epistemic/practical/moral/whatever else divide.

As I hope mundane cases like DANCE and SMOKE make clear, the intra-contextual variability of reasons-claims is extremely widespread. For those who take our ordinary judgments involving the count

noun 'reason(s)' at face value, however, such variability should be disconcerting. For it shows that ordinary reasons-claims are far more sensitive to conversational context than previously realized. In particular, it shows that what can legitimately count as a reason (or among the reasons) to do something in a particular context—just like what can legitimately count as "the" reason—depends in part on idiosyncratic facts about *us*, including what else has been cited as a reason.[199] But that's not what we would expect if reasons-claims corresponded directly to the underlying normative facts, which presumably do *not* so depend. To put the point slightly differently: if (true) claims we make using sentences of the form '*r* is a reason (for *S*) to *φ*' expressed facts that were wholly objective (or 'out there in the world'), then a given fact *r*'s putative status *as a reason* in a given context presumably wouldn't depend on obviously irrelevant factors, such as those having to do with conversational dynamics. Instead, we would expect the infelicity of any such claims to be *merely* pragmatic—to be pointless, irrelevant, misleading in virtue of generating a false implicature, or some such, despite being strictly speaking true. But the defectiveness of (D3)-(D5) and (S3)-(S5) is more than merely pragmatic, and isn't traceable to one of the usual suspects.[200] To insist that they're strictly speaking true but suffer from some familiar pragmatic malady doesn't do justice to the character of their defectiveness, which is stubbornly present in a way that familiar pragmatic maladies are not.

It's important to be clear that I'm *not* claiming it's impossible for (D3)-(D5) and (S3)-(S5) to be used felicitously—on the contrary, they can be. Suppose, for instance, that we discover that Billy gets

---

[199] Again, this parallels Swanson's (2010) point about the contextual variability of our judgments involving 'a cause', extending Lewis's (1973) point with respect to 'the cause'.

[200] The "usual suspects" qualifier is included because on the view I sketch below, the defectiveness can be seen as straddling the semantic/pragmatic boundary, and so might be considered at least quasi-pragmatic in nature. In this respect the defectiveness of (3)-(5) might be compared—but not assimilated—to the defectiveness involved in cases of presupposition failure, which similarly straddles the semantic/pragmatic line. What's important is just that the defectiveness is not pragmatic in any of the familiar or straightforward ways that we would expect were reasons-claims reliable guides to the underlying facts.

paid handsomely for every dance party he attends. In that context (D3)-(D5) are fine. So the infelicity doesn't attach to the sentences themselves, but rather to their use in the context. For this reason, among others, the account I'll offer of ordinary reasons-claims won't be directly concerned with the *truth conditions* of such claims. Instead, the primary focus will be on their *acceptability* conditions—i.e. the circumstances in which it is acceptable for someone who knows all the relevant non-semantic facts to use sentences of the form '*r* is a reason for *S* to φ'—and how those conditions are sensitive to context.[201]

I'll provide an explanation of why (D3)-(D5) and (S3)-(S5) are acceptable in some contexts and not in others below, but I mention it now just to make clear the scope of my claim: I'm only claiming that (D3)-(D5) and (S3)-(S5) can't be used felicitously *given knowledge of the relevant non-semantic facts in* DANCE and SMOKE. This shows why the diagnosis that (D3)-(D5) and (S3)-(S5) are merely guilty of implicating something false is unsatisfactory—implicatures are characteristically cancelable (a la Grice (1989)), but I'm not aware of any continuations of (D3)-(D5) and (S3)-(S5) in the relevant contexts that eliminate the infelicity. As we'll see, a semantic explanation—one that concerns the communicative role of reasons-talk—is more promising.

## 7.3 Reasons as representatives

The immediate moral to be drawn from the intra-contextual variability of reasons-talk is that insofar as we're interested in carving normative reality at its joints, or at least as close to its joints as we can, we should be wary of putting much stock in our judgments concerning what and whether

---

[201] Cf. Swanson (2010). If a truth-conditional account is desired, however, it can be provided. For example, one might relativize sentences of the form '*r* is a reason for *S* to *φ*' to conversational contexts, where such contexts are stipulated to include information about (e.g.) what else has been cited as a reason. See Strevens (2013) for a roughly analogous account of the truth conditions of sentences of the form '*c* is a cause of *e*', which he relativizes to "explanatory frameworks".

something is (i.e. *felicitously counts as*) a reason for something else. More generally, however, I think the lesson to be learned is that we shouldn't take singular explanations—i.e. explanations where the explanans is singular—of normative facts at face value, whether the normative facts to be explained are

**contributory** (e.g. *f is a reason for S to (not) φ; f gives S reason to (not) φ*[202])

or

**threshold-y** (e.g. *f makes it good/bad/(un)reasonable for S to φ; f is why S may/must/should (not) φ*)

Given what we know about the nature of our explanatory practices in general, in which surprisingly few facts usually suffice (for communicative purposes, *not* metaphysically) to explain complicated facts or events, it should come as little surprise that when we make claims of the latter "threshold-y" form the cited fact *alone* will rarely be responsible for the relevant normative status of *S*'s φ-ing, whether it be good, bad, reasonable, unreasonable, or whatever. Instead, other facts will almost always play an important role as well, with full responsibility for the normative status of *S*'s φ-ing belonging to a cluster of facts that includes but is not limited to the fact cited. Call any cluster of facts that taken together—and only taken together—are normatively relevant (in some particular way, to some degree) to some action or attitude a "normative cluster". For reasons having to do with communicative efficiency as well as our informational limitations, we rarely mention entire normative clusters. Instead, we single out one or two particularly relevant and/or accessible facts and ascribe responsibility to them. In essence, the cited facts function communicatively as representatives of, or proxies for, the normative clusters to

[202] Judgments of the form '*f* gives *S* reason to φ' are subject to the same kind of contextual variability as those of the form '*r* is a reason for *S* to φ'. They are both singular explanations of normative statuses—namely, that there is (a certain amount of) reason for *S* to φ—and so this is to be expected. Rarely do singular facts do all the relevant explanatory work on their own.

which they belong, or are otherwise related.[203] On the resulting view, a fact playing such a role should not (by itself) be taken as indicative of anything metaphysically or explanatorily deep. In a different conversational context—one differing in terms of shared background information, preceding discourse, etc.—some other fact belonging to the very same cluster that fully explains the relevant normative status of *S*'s *φ*-ing might be more salient, or otherwise relevant, and hence be referenced instead.[204]

What's more surprising—or at least what's been insufficiently remarked upon—is that precisely the same thing is true when it comes to "contributory" claims of the form '*f* is a reason for *S* to φ', which (as I've argued) are essentially explanatory in nature. Accordingly, when making or evaluating such claims we shouldn't take a fact's status *as a reason* (or "reason-giver") to be indicative of anything metaphysically or explanatorily deep. For as we've seen, whether a given fact can felicitously count as a reason in a particular context typically depends not only on there being other facts which taken together —and only taken together—fully explain the relevant normative status of an agent's action or attitude, but *also* on whether any of those other facts have already been cited as a reason. This helps explain why a dispute over whether we should accept (D1) or instead (D2) in the context of DANCE, or (S1) or instead (S2) in SMOKE, will (according to many people) seem shallow. Each fact is eligible to serve as a representative of the relevant normative cluster, and it's up to us whether we use one or the other (or their conjunction). The facts we cite as reasons merely function as representatives; they're not themselves the real thing.

---

[203] There is an obvious analogy with Mackie's (1965) INUS account of causation (or causal explanation), though there are also some important differences. See Strevens (2013) for a recent defense of a (non-reductive) Mackie-style account of causal claims that is congenial to the view of reasons-claims that I develop.

[204] See Broome (2013)—especially pp. 48-51—for similar remarks concerning our explanatory practice when it comes to (what I'm calling) threshold-y claims.

Understanding reference to reasons in this way—i.e. as reference to particular facts that function as (something like) representatives of the normative clusters to which they belong, or are otherwise related—allows for a tidy explanation of the otherwise problematic contextual variability observed in cases like DANCE and SMOKE. For on this view it's natural to take our thought and talk about reasons to be implicitly governed by something like the following norm:[205]

> **Use Good Representatives**
>
> When you attribute some normative support for A to a normative cluster relevant to A, use good representatives of that cluster for the purposes at hand.

The relevant sense of 'goodness' is a broadly communicative one. The goodness of a representative is function of the amount of information it provides—either directly or indirectly, given the background information—about the normative cluster to which it belongs, as well as the economy with which it provides that information. Ceteris peribus, the better a representative satisfies this dual constraint, the better a representative it is. Typically the most apt representatives of a given normative cluster will be those that are the biggest "difference makers" in the sense that changes in them, holding other features of the situation fixed, would make the most difference to the normative status of the subject's attitude or action.

Unsurprisingly, the relative "goodness" of a potential representative will vary from context to context, and even within a context, depending on the interests of and background information shared by the conversational participants—including whether (and if so, which) other representatives have already

---

[205] This principle intentionally mirrors (in name and form) one that Swanson (2010) uses to explain the contextual variability of our causal-talk. Once again, much of my discussion in this section intentionally mirrors Swanson's discussion of causation in order to illustrate the similarities between our thought and talk about reasons and causes.

been used. Consider again DANCE, where—let's suppose—the normative cluster consists of just two facts (I'll use '[*p*]' as a noun phrase that refers to the fact that *p*):

(**NC**) {[Billy loves to dance]; [There will be dancing at the party]}

As our judgments concerning (D1)-(D2) demonstrate, the initial conversational context is one in which either fact belonging to (NC) is eligible to serve as a good representative:

(**D1**) The fact that there will be dancing at the party is a reason for Billy to attend.

(**D2**) The fact that Billy loves to dance is a reason for him to attend.

But although both facts start off on a communicative par, the act of citing either as a reason robs the other of its potential role as a representative. In this way, using the fact that there will be dancing to represent (NC) renders impotent—mid-utterance—the representative role that Billy's love of dancing could have otherwise played. In the same way, using the fact that Billy loves to dance to represent (NC) renders impotent—mid-utterance—the representative role that the fact that there will be dancing could have otherwise played. To nonetheless cite both facts individually as reasons would be to add unnecessary (and indeed misleading) information.[206] Hence the infelicity of (D3):

(**D3**) #The fact that there will be dancing at the party is a reason for Billy to attend, and so is the fact that he loves to dance.

---

[206] Notice that despite being overkill it would be fine to cite both facts collectively:

(D1+2) The fact that there will be dancing at the party together with the fact that Billy loves to dance is a reason for him to attend.

The same reasoning applies to (D4)-(D5), and an analogous explanation is available for (S3)-(S5) in SMOKE.

Use Good Representatives also helps explain why the default interpretations of (D3)-(D5) are ones where it's assumed that Billy's love of dancing and the fact that there will be dancing at the party are normatively independent of each other.[207] Since both facts are being used as representatives, it will be assumed that both are *needed*—that both are playing an important representative role. But the only way for each of them to play such a role would be for them to represent distinct normative clusters. This is why (D3)-(D5) are fine in a context where (e.g.) Billy gets paid handsomely for attending dance parties. The same point applies to (S3)-(S5) in SMOKE)

## 7.4 Two alternatives: narrowing and coarsening

What the foregoing suggests is that there's a significant amount of distance between our ordinary judgments involving the count noun 'reason(s)' and the underlying normative facts. This conclusion isn't forced on us, however, and those who wish to maintain a closer connection between the two might try to provide an alternative account of the intra-contextual variability data. There are two main lines of response that I'll consider. To streamline the discussion, I'll focus on how each response handles (D1)-(D5) in DANCE, although everything I say applies, mutatis mutandis, to (S1)-(S5) in SMOKE, as well as countless other examples.

According to the first line of response, only one of (D1) or (D2) truly specifies a reason for Billy to attend. The other (at best) merely specifies a *background condition* for the reason—i.e. a condition that needs to obtain in order for the other fact to constitute a reason, but which is not itself part of the

[207] "Normatively independent" in the sense that the normative relevance of at least one of the facts doesn't depend on the other for its belonging to a normative cluster relevant to Billy's going to the party.

reason. On this view, then, only one sentence is true, and the acceptability of the other is to be explained on other grounds. Since this attempts to preserve a finely-individuated picture of reasons at the cost of reinterpreting or rejecting various ordinary reasons-claims, I'll call this the "narrowing" response.

The second line of response is to argue that neither (D1) nor (D2) identify the "real" or "full" reason for Billy to attend (at least not directly). Instead, they merely specify a *part* of the reason, with context providing the rest of the information needed to identify the reason itself—i.e. the conjunctive fact that Billy loves to dance and there will be dancing at the party. This nicely explains why (D3)-(D5) seem to involve double-counting, and why it would strike us as a mistake to list the two facts separately under the "pro" column while deliberating. Because it individuates reasons more coarsely than (many of) our ordinary reasons-claims do, I'll call this the "coarsening" response.

Neither response is fully satisfying. There's a lot that could be said, but I'll focus on just a couple points of dissatisfaction. Consider first the coarsening response. Assuming the proponent of this response purports to be using *reason(s)* in its ordinary normative sense, she's going to be committed to thinking that nearly all of our reasons-claims, including (D1) and (D2), are false. The fact that there will be dancing at the party, it turns out, is *not* a reason for Billy to attend, and neither is the fact that he likes to dance. Neither fact is a reason any more than half of a car is a car. Although the bitterness of an error theory might be sweetened by supplementing it with a plausible pragmatic story, it's still something that many would rather avoid. The most obvious way of avoiding an error theory is to say—as I do—that what counts as a reason is highly context-sensitive matter. But since the underlying normative facts are presumably *not* a highly context-sensitive matter (at least not in the same way), this is to concede the point that our reasons-claims fail to directly correspond to the underlying normative facts.

Accordingly, the proponent of the coarsening response might shift tactics and try to distance herself from the ordinary sense of 'reason(s)' by introducing a technical term—most creatively, she could use 'Reason(s)'—to better capture what our ordinary reasons claims only capture in part.[208] The conjunctive fact, but neither of its components, would then be (e.g.) a Reason for Billy to attend, with (D1) and (D2) merely being used to direct our attention to its vicinity by identifying a part of it. There are various ways in which this strategy might be fleshed out, and the account of reasons-claims offered above can be understood as one way of doing so. The main point is just that such a response acknowledges the distance between ordinary reasons-claims and the underlying normative facts, and it's a substantive further issue just how close or far apart they are. Such claims would therefore need to be treated with caution and rarely, if ever, taken at face-value. Methodologically, it would be illegitimate to rely on ordinary reasons-claims in motivating (or attacking) a certain principle or theory while formulating the official principle or theory in other terms, without a careful story connecting the two.

What about the narrowing approach? Once again a lot could be said, but it's worth emphasizing just how widespread the intra-contextual variability of reasons-talk is. For if the narrowing response is adopted across the board, we'll end up having to privilege a large range of intuitively acceptable reasons-claims at the expense of others, and in each case we'll need to both motivate the purported asymmetry as well as explain why it's something we're insensitive to. In some cases the motivation and explanation might seem easy to provide. For instance, various philosophers have denied that desires are themselves

---

[208] Raz (2000), for instance, introduces the term 'complete reason' for stand for "all the facts stated by the non-redundant premises of a sound, deductive argument entailing as its conclusion a proposition of the form 'There is a reason for *P* to *V*'". The utility of such a notion is questionable—see Dancy (2004) for criticism. Scanlon's (2014) introduction of the normatively fundamental four-place 'R' relation is a more promising move along such lines, though as with Raz there's a tension between the introduction of a technical term and the often rather casual reliance on ordinary reasons-claims. Indeed, the intra-contextual variability data strongly suggests that Scanlon's R-relation either fails to carve normative reality at its joints, or else is significantly more removed from ordinary reasons-claims than he realizes.

reason-giving, and so would reject (D2) in DANCE.[209] Why do so many of us find (D2) acceptable, then? That'll depend. Perhaps we're just mistaken in thinking that desires are reason-giving, and so (D2)'s veneer of plausibility is simply the result of a widespread but nonetheless erroneous assumption.[210] Alternatively, perhaps desires are merely necessary background conditions for reasons and not (part of) reasons themselves.[211] In that case (D2)'s plausibility might be explained along more pragmatic lines—even though Billy's love of dancing isn't a reason, it's closely related to something that is.

It's worth noting, however, that the explanatory debts don't end there. For it's not enough to explain why, say, (D2) in DANCE seems true even though it's not; one also has to explain why it *doesn't seem false*. That is, we have to explain why in the context of DANCE neither of these are acceptable:

(**D2***) The fact that Billy likes to dance is not a reason for him to go.

(**D3***) The fact that there will be dancing at the party is a reason for Billy to go, but the fact that he likes to dance is not a reason for him to go.

This additional explanatory burden applies to advocates of the coarsening response just as much as it does to advocates of the narrowing response. And it's one that the account of reasons as representatives is well-placed to discharge.

Let's suppose, though, that the narrowing approach is plausible in the case of DANCE. That does does little to motivate it in general, since the intra-contextual variability of reasons-talk extends far beyond cases involving desires. Consider SMOKE, for instance. What grounds do we have for privileging

---

[209] Or at least provide a non-standard explanation of why it's true.

[210] Cf. Scanlon (1998), Dancy (2004).

[211] Cf. Schroeder (2007).

one of (S1) or (S2) over the other? When it comes to believing there's a fire, is the sight of smoke epistemically relevant in a way that the known connection between smoke and fire is not? Here it seems even less plausible to insist that only one of (S1) or (S2) specifies the "real" reason, while the other merely specifies a background condition. The coarsening response has more plausibility here. The defender of a close relationship between ordinary reasons-claims and the underlying normative facts might therefore opt for a mixed strategy, treating each case on its own terms.

No matter which approach is adopted, however, we'll have to concede that our ordinary judgments about reasons do not reliably track the fundamental normative facts, and we'll have to explain why we're systematically inclined to cite as reasons facts that aren't in fact reasons. Although I'm not in a position to rule out the possibility that such a story can be told that applies to the full range of cases, I'm also not optimistic. And other things being equal, we should prefer a non-error-theoretic account.

## 7.5 Some Consequences

So far I've considered two different challenges facing those who place a lot of theoretical weight on our ordinary normative judgments involving the count noun 'reason(s)', the first arising from the relationship between the count noun and the mass noun and the second arising from the intra-contextual variability of reasons-claims. It's worth emphasizing, however, that the mere fact that a theory or principle is formulated primarily using the count noun 'reason(s)'—in its ordinary normative sense—isn't always a problem. And that's because it will usually be possible to reformulate the theory or principle in such a way that it no longer makes use of the count noun but is in all important respects equivalent to the original. But other times it won't be possible to reformulate the theory or principle in this way, and it's in these cases (and only these cases) that I think the theories or principles should be

viewed with suspicion. The recent emphasis and burgeoning literature on reasons in (meta)ethics and epistemology furnishes us with a myriad of cautionary tales, but for the purposes of illustration I'll just consider two. In particular, I'll explain how the foregoing considerations threaten to undermine the main argument for moral particularism as well as Schroeder's recent defense of a Humean theory of practical reasons.

## 7.6 Against particularism (moral or otherwise)

Particularists—as I'll be understanding them—about a given normative domain *D* (e.g. moral, practical, epistemic) are concerned to deny that there are any substantive, finite, and exception-less principles concerning *D*.[212] As Berker (2007: 112) explains, in arguing for their view particularists have historically assumed

> ...a three-level framework according to which non-normative facts at the first level determine the facts about reasons at the second level, which in turn determine other normative facts such as the overall rightness and wrongness of actions at the third level, and they [claim] that there is variability of an inscrutably complex sort both in how the second level depends on the first and in how the third level depends on the second.

We can thus can understand particularism as a two-fold thesis: it concerns both the way in which reasons arise out of a situation's non-normative features and the way in which the reasons that are

[212] A bewildering variety of theses have fallen under the label "particularism" over the past couple decades. I merely intend to capture one particularly prominent strain of particularist thought. A proper examination of the consequences for the many different particularist theses would require a much more exhaustive treatment.

present in a given situation combine to yield the overall normative status of a possible action or attitude (call it $\varphi$).[213] More fully:

> **holism about reasons:** for every non-normative feature that is a reason for/against $\varphi$-ing in one possible context, there is another possible context in which that same feature either is a reason of opposite valence or else fails to be a reason of any kind.[214]
>
> **noncombinatorialism about reasons**: the combinatorial function takes as input the valence and weight of all the reasons in a given possible situation and gives as output the normative status (e.g. rightness or wrongness) of each $\varphi$ is not finitely expressible (and so not additive).

Although the argument for particularism based on holism has been challenged on a number of fronts (see especially Ridge and McKeever (2006)), whatever force it retains is undermined once we realize that reasons are merely representatives—that is, once we understand reference to reasons as reference to facts that serve as representatives of normative clusters. Since particular facts may belong to, and hence serve as representatives of, different normative clusters in different (non-conversational) contexts, it's only to be expected that, as Margaret Little (2001: 34) puts it, a "consideration that in one context counts [as a reason] for an action, can in another count against it or be irrelevant". Such contextual variability is straightforwardly accommodated by the account of reasons as representatives—an account

---

[213] This, along with what immediately follows, is a generalization of Berker's characterization of moral particularists on pp. 119-122.

[214] There are weaker formulations of holism that are not committed to the universal generalization being true —as both Dancy (2004) and McKeever and Ridge (2006) note—though exceptions will be treated as "special".

which is intentionally silent about the nature of the underlying normative facts—and so it doesn't, by itself, have any robust normative or metaphysical implications. In this way, the argument from holism (at least in its usual guise) is effectively neutered, since it wrongly takes the contextual variability of reasons-claims to reveal something significant about the normative realm when it's really just an artifact of our reasons-talk. Non-particularists can thus accommodate, and indeed predict, the relevant data without conceding the metaphysics. Support for particularism will have to come from elsewhere—the variability of our judgments about reasons is a non-starter.

Importantly, one might still deny that there are any substantive, finite, and exception-less principles in a given normative domain *D* by arguing that the normative forces (support, opposition, etc.) generated by *D*-clusters do not interact in a way that is finitely systematizable. (This concerns the link between the second and third levels above, rather than the first and second.) But this would only seem to vindicate a kind of normative *regionalism*, not holism, since the normative relevance of (at least some) normative clusters is arguably constant across contexts, even if the interactions between clusters that together determine all-things-considered statuses remain invariably complex.

## 7.7 Against Hypotheticalism

Schroeder (2007) ably and imaginatively defends a (neo-)Humean theory of reasons that he calls Hypotheticalism. At the heart of this view is the following claim:

> **Reason*** For all propositions *r*, agents *x*, and actions *a*, if *r* is a reason for *x* to do *a*, that is because there is some *p* such that *x* has a desire whose object is *p*, and the truth of *r* is part of what explains why *x*'s doing *a* promotes *p*. (29)

If true, **Reason*** would provide what many philosophers would love to have: a reduction of normative reasons (at least of the practical variety), where facts about reasons are explained in terms of, and ultimately grounded in, facts about the psychological states (in the case of Humean theories, desires) of agents. We already know that in some cases some reasons are to be explained in terms of the agent's psychology, and it would be nice if we could explain all reasons in such terms.

Of course, Humean theories face a familiar battery of objections, and much of Hypotheticalism's attractiveness is due to its apparent ability to avoid or blunt many of them. In particular, central to Schroeder's defense of Hypotheticalism is the claim that there is a principled distinction to be drawn between *reasons* and *background conditions*, and that desires merely belong to the latter. Now, the distinction between a kind of thing and what explains why something is a thing of that kind is, in general, a perfectly good distinction. But Schroeder realizes there are certain cases in which such a distinction has "proved to be hard to maintain—e.g. in the literature on causation" (24). This is important because of the striking and systematic parallels between our talk of causes and our talk of reasons—parallels that I've only been able to hint at so far.

For example, that just as there are singular ascriptions, both specific and non-specific, of reasons, so there are singular ascriptions of causes:

**Singular ascriptions**

(R1) The fact that there will be dancing at the party is a reason for Billy to go.

(C1) Sally's sleep apnea is a cause of her fatigue.

(R2) There is a reason for Billy to go to the party.

(C2) There is a cause of Sally's fatigue.

Similarly, just as there are general (or generic) ascriptions of reasons, so there are general ascriptions of causes:

**General ascriptions**

(R3) The fact that there will be dancing at a party is a reason to go.

(C3) Sleep apnea is a cause of fatigue.

Moreover, just as there are related "threshold-y" and "contributory" normative locutions, so there are "threshold-y" and "contributory" causal locutions:

**Threshold-y ascriptions**

(R4) Billy should go because there will be dancing at the party.

(C4) Sally is fatigued because of her sleep apnea.

**Contributory ascriptions**

(R5) The fact that there will be dancing is relevant to whether Billy should go.

(C5) Sleep apnea contributes to Sally's fatigue.

In addition to the fact that 'reason(s)' and 'cause(s)' pattern linguistically in very similar ways, they also both exhibit the same kind of intra-contextual variability explored above, and they are both plausibly governed by Use Good Representative-like norms. (For a defense of the latter two claims in the case of causation, see Swanson (2010).) Such similarities become less surprising once it's realized that reasons

and causes are both "becauses"—i.e. they are things (facts, states, events, objects, whatever) that help explain things.

Given the systematic similarities between our reasons-talk and causal-talk, the challenge for Schroeder is straightforward: why think the needed distinction between reasons (in the ordinary normative sense) and background conditions is any more metaphysically significant than the distinction between causes and background conditions? Consider the standard example of a match being struck and catching on fire. Although some have argued that only the striking of the match is a (or "the") cause of the fire, with the presence of oxygen merely being a background condition, it's widely held that such discrimination between causally relevant factors is untenable and shouldn't be reflected in our metaphysics. For there are other contexts in which it's clear that the underlying metaphysics remains the same and yet what we're inclined to count as a (or "the") cause changes. (Suppose, for example, the match is being repeatedly struck in a vacuum when oxygen suddenly rushes in.)

Why not think the same lesson applies in the case of reasons? That is, just as (e.g.) the striking of a match and the presence of oxygen are on a causal par—they're both part of the same causal process terminating in the presence of fire and the causal relevance of each depends on the presence of the other —why not treat the fact that (e.g.) Billy loves to dance and the fact that there is dancing at a party as being on a normative par? Why not think that they're both part of the same normative cluster bearing on whether Billy should go and the normative relevance of the one depends on the presence of the other (together forming what I've called a "normative cluster")? In the former case we have good reason to think the distinction between causes and background conditions isn't metaphysically robust, and that it tells us more about our interests and information than reality. Given the parallels between reasons and causes, the needed distinction between reasons and background conditions appears to fare no better.

## 7.8 The "flattening" worry

One important worry with the account I've explored so far is that it seems to ignore important distinctions between different explanatory roles that the facts belonging to a given normative cluster might play. Denying the distinction between reasons and background conditions makes this worry vivid. For recall that normative clusters are just clusters of facts that, taken together, are normatively relevant (in some particular way, to some degree) to some action or attitude. To take normative clusters as explanatorily basic—in the metaphysical sense—seems tantamount to "flattening" the explanatory landscape, treating all facts within them as being on a par. But this might seem problematic—we might want to distinguish between different (metaphysically robust) explanatory roles that facts belonging to a given normative cluster might play, such as (cf. Dancy (2004), Bader (forthcoming)):

(a) that which *makes it the case* that there is normative support for $\varphi$ (Sources)

(b) that which (merely) *enables* there to be such support, and (Enablers)

(c) that which (merely) *modifies* the amount of support. (Modifiers)

By ascribing ultimate explanatory responsibility to normative clusters rather than the facts cited as reasons, aren't I guilty of flattening the explanatory landscape?

No, I'm not. Everything I've said is compatible with the view that there are different metaphysically robust explanatory roles to be played by different (sub-clusters of) facts. Even if full explanatory responsibility resides with a given cluster, that's compatible with different facts within it playing different roles, whether on their own or in conjunction with others. Indeed, the distinctions above are just particular applications of more general metaphysical distinctions that might, with equal or

more justice, be applied to a wide range of normative notions, such as value, evidence, justification, rights, obligations, and, of course, reason (mass).

In fact, similar distinctions can be (and are) drawn outside the normative realm altogether. For instance, we might want to distinguish the "grounds" of friendship—i.e. what it is in virtue of which someone is a friend—from the mere "conditions" of friendship, as well as from various possible "modifiers" of one's status or standing as a friend. Or we might not. Whether such distinctions mark genuine differences between explanatory roles is something to be decided on a case-by-case basis. The story I've told about reason(s)-claims and their relationship to normative clusters is largely silent on the extent to which the such distinctions apply in this particular case, though the intra-contextual variability data does suggest that they normally fail to apply directly to the facts we cite as reasons. At best the distinctions apply to larger clusters of facts that include the ones we cite.

In fact, the promiscuity of reasons-claims suggests something stronger—namely, that the facts we regularly and unproblematically cite as reasons fail to play the same metaphysical (as opposed to *communicative*) role, even when taken together with other facts.[215] For examples, a lot of philosophers want to deny that the ability to $\varphi$ is itself a reason—or part of a reason—to $\varphi$, but there are many contexts in which it seems perfectly fine to cite one's ability to $\varphi$ as a reason to $\varphi$. These include contexts in which $\varphi$ is obviously choiceworthy (e.g. 'The fact that you can help someone in need is a reason to do so') as well as those in which one isn't usually able to $\varphi$ (e.g. 'It's rare to get the chance to ask the

[215] Uncontroversial examples are hard to come by; what follows is for illustrative purposes.

President a question. The fact that you're able to right now is a reason to do so.').[216] We face a similar range of choices with respect to such claims as those in DANCE and SMOKE. One option is to bite the bullet and say that they're strictly speaking false, offering some pragmatic story to explain (or explain away) their acceptability. Another—which I favor—is to accept them as perfectly good reasons-claims, since the facts cited are good representatives in the imagined contexts and that's all that being (i.e. counting as) a reason really amounts to. Nothing would then follow about the specific metaphysical role the facts play within the cluster they represent—that would be a further, substantive issue, to be settled on broadly theoretical grounds.

Regardless of which option we choose, however, the point remains that ordinary reasons-claims tell us very little on their own. All they tell us is that the fact cited, together with other facts, plays some role in explaining—in the pragmatic, but not necessarily metaphysical, sense—why there is reason to perform the relevant action or adopt the relevant attitude. They don't tell us *which* other facts those are, *what* the specific explanatory role is, *how much* reason there is, nor what *kind* of reason there is. All that a claim of the form '*r* is a reason (for *S*) to *φ*' entails is that there's something (not: some *thing*) to be said in favor of *φ*-ing, and that *r* is something it makes sense to cite in an effort to explain why. We need to rely on background information and engage in further theorizing to fill in all the gaps.

## 7.9 Conclusion

---

[216] The same applies to the following examples provided by Dancy (2004):

1. I promised to do it.
2. My promise was not given under duress.
3. I am able to do it.

A. She is in trouble and needs help.
B. I am the only other person around.

Dancy takes (1)/(A) to be Sources while (2)-(3)/(B) are Enablers/Modifiers, but he doesn't grapple with the fact that each can be felicitously cited as a reason in normal contexts, even when the normatively relevant considerations are held fixed.

I've explore various subtleties in how we ordinarily think and talk about reasons, revealing some of the limitations and liabilities of taking such talk seriously. Somewhat ironically, the challenges I've considered—those arising from the use of *reason* as both a count noun and as a mass noun as well as from the context-sensitivity of reasons-claims—are the result of taking the language of reasons *more* seriously than it has been taken by many, if not all, of its proponents. It turns out that our judgments about reasons are, at best, a partial and highly defeasible guide to what really matters, both normatively and metaphysically.

# BIBLIOGRAPHY


Bader, Ralf (forthcoming). Conditions, Modifiers, and Holism. In *Weighing Reasons.* Oxford University Press.

Beall, Jc (2009). *Spandrels of Truth*. Oxford University Press.

Bennett, Jonathan (2003). *A Philosophical Guide to Conditionals*. Oxford University Press.

Bergmann, Michael (2005). Defeaters and Higher-Level Requirements. *The Philosophical Quarterly, 55,* 419-36.

Berker, Selim (2007). Particular Reasons. *Ethics*, *118*(1): 109–139.

—— (ms) 'Graphic Coherence'.

Broome, John (1999). Normative Requirements. *Ratio*, *12*(4), 398–419.

—— (2004). Reasons. In *Reason and Value: Essays on the Moral Philosophy of Joseph Raz*. 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 (2013). Reasons as Explanations. *Philosophical Studies*, *165*(3): 805–24.

Caie, Michael (2013). Rational Probabilistic Incoherence. *Philosophical Review, 122*(4), 527–75.

Carr, Jennifer (ms). Imprecise Evidence without Imprecise Credences .

Christensen, David (2004). *Putting Logic in its Place.* Oxford University Press.

—— (2008). Does Murphy's Law Apply in Epistemology? *Oxford Studies in Epistemology, Vol. 2*.

—— (2010) Higher-Order Evidence. *Philosophy and Phenomenological Research*, *81*(1), 185–215.

Dancy, Jonathan (1977). The Logical Conscience. *Analysis, 37*(2), 81–84.

—— (2004). *Ethics Without Principles*. Oxford University Press.

Darwall, Stephen (1983). *Impartial Reason.* Cornell University Press.

Dowell, Janice (2011). A Flexible Contextualist Account of Epistemic Modals. *Philosophers' Imprint, 11*(14).

—— (2012). Contextualist Solutions to Three Puzzles about Practical Conditionals. In *Oxford Studies in Metaethics*, *Vol. 7*.

—— (2013). Flexible Contextualism about Deontic Modals. *Inquiry, 56*(2-3), 149–178.

Dowell, Janice, and Bronfman, Aaron. (forthcoming-a). Contextualism about Deontic Conditionals. In *Deontic Modality*. Oxford University Press.

—— (forthcoming-b). The Language of Reasons and 'Ought's. In *The Oxford Handbook of Reasons and Normativity*. Oxford University Press.

Easwaren, Kenny, and Branden Fitelson (forthcoming). Accuracy, Coherence, and Evidence. In *Oxford Studies in Epistemology, Vol. 5*.

Edgington, Dorothy (1995). On Conditionals. *Mind, 104*.

—— (2006). Conditionals. In *The Stanford Encyclopedia of Philosophy*.

Enoch, David (2011). On Mark Schroeder's Hypotheticalism. *Philosophical Review, 120*(3): 423-446.

Field, Hartry (2009). Epistemology without Metaphysics. *Philosophical Studies 143*(2),249–90.

Finlay, Stephen (2014). *Confusion of Tongues: A Theory of Normative Language*. Oxford University Press.

Friedman, Jane (2013a). Suspended judgment. *Philosophical Studies, 162*(2), 165–181.

—— (2013b). Rational agnosticism and degrees of belief. In *Oxford Studies in Epistemology, Vol. 4*, 57–82.

Gensler, Harry J. (1985). Ethical Consistency Principles. *The Philosophical Quarterly 35*(139), 156–70.

Gibbons, John (2010). Things That Make Things Reasonable. *Philosophy and Phenomenological Research, 81*(2), 335–361.

Gillies, Anthony S. (2010). Iffiness. *Semantics and Pragmatics, 3*(4), 1–42.

Gillon, Brandon S. (2012). Mass Terms. *Philosophy Compass, 7*(10): 712–730.

Greco, Daniel. (2014). A Puzzle about Epistemic Akrasia. *Philosophical Studies 167*(2): 201–19.

Greenspan, Patricia (1975). Conditional Oughts and Hypothetical Imperatives. *Journal of Philosophy, 72*(10), 259–76.

Grice, H. Paul (1989). *Studies in the Way of Words*. Harvard University Press.

Hacquard, Valentine (2010). On the event relativity of modal auxiliaries. *Natural Language Semantics, 18*(1), 79–114.

Harman, Gilbert (1986). *Change in View*. MIT Press.

—— (2009). Field on the Normative Role of Logic. *Proceedings of the Aristotelian Society, 109,* 333-335.

Hill, Thomas E. (1973). The Hypothetical Imperative. *The Philosophical Review, 82*(4), 429–450.

Holton, Richard (2008). Partial Belief, Partial Intention. *Mind, 117*(465), 27–58.

Huddleston, Rodney, and Pullum, Geoffrey. K. (2002). The Cambridge Grammar of the English Language. Cambridge University Press.

Jackson, Frank (1998) *From Metaphysics to Ethics*. Clarendon Press.

Joyce, James M. (2005). How Degrees of Belief Reflect Evidence. *Philosophical Perspectives, 19,* 153-179.

—— (2011). A Defense of Imprecise Credences. *Oxford Studies in Epistemology, Vol. 4*.

Kelly, Thomas (2014). Evidence Can Be Permissive. In *Contemporary Debates in Epistemology* (2nd ed., 298-313). Blackwell.

Khoo, Justin. (2011). Operators or Restrictors? A Reply to Gillies. *Semantics and Pragmatics*. Volume 4, Article 4: 1-25.

—— (2013). Conditionals: Interactions with Adverbs, Modals and Tense. (Dissertation)

Kim, Jaegwon (1998). *Mind in a Physical World*. MIT Press.

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

—— (2007). State or Process Requirements? *Mind, 116*(462), 371–385.

—— (2007). How Does Coherence Matter? *Proceedings of the Aristotelian Society, 107*.

—— (2008). The Myth of Practical Consistency. *European Journal of Philosophy, 16*(3), 366–402.

Kratzer, Angelika (2012). *Modals and Conditionals*. Oxford University Press.

Lance, Mark, and Margaret Little (2004). Defeasibility And The Normative Grasp Of Context. *Erkenntnis 61*(no. 2–3): 435–55.

Lewis, David K. (1973). Causation. *Journal of Philosophy,* 70: 556–567.

Little, Margaret O. (2001). On Knowing the "Why": Particularism and Moral Theory. *The Hastings Center Report, 31*(4): 32–40.

Lord, Errol (2014a). The Real Symmetry Problem(s) for Wide-Scope Accounts of Rationality. *Philosophical Studies, 170*(3),443-464.

—— (2014b). The Coherent and the Rational. *Analytic Philosophy*.

Mackie, John L. (1965). Causes and Conditions. *American Philosophical Quarterly, 2*(4).

McKeever, Sean D., and Ridge Michael R. (2006). *Principled Ethics: Generalism as a Regulative Ideal*. Oxford University Press.

Nickel, Bernhard (forthcoming). *Between Logic and the World: An Integrated Theory of Generics*. Oxford University Press.

Olson, Jonas. (2009). Reasons and the New Non-Naturalism. In *Spheres of Reason: New Essays in the Philosophy of Normativity*. Oxford University Press.

Parfit, Derek. (2011). *On What Matters*, Vol. I. Oxford University Press.

Pietroski, Paul (1993). Prima Facie Obligations, Ceteris Paribus Laws in Moral Theory. *Ethics 103*(3): 489–515.

Priest, Graham (1979). The Logic of Paradox. *Journal of Philosophical Logic*, 8, 219–241.

—— (2006). *In Contradiction* (2nd ed.). Oxford University Press.

Pryor, James (2004). What's Wrong with Moore's Argument? *Philosophical Issues*, *14*(1), 349–378.

Quinn, Warren S. (1990). The Puzzle of the Self-Torturer. *Philosophical Studies*, *59*(1), 79–90.

Rawls, John (1971). *A Theory of Justice*. Harvard University Press.

Raz, Joseph(2000). Engaging Reason. Oxford University Press.

—— (2005). The Myth of Instrumental Rationality. *Journal of Ethics and Social Philosophy*.

Rothschild, Daniel (forthcoming). A Note on Conditionals and Restrictors.

Scanlon, Thomas M. (1998). *What We Owe to Each Other*. Harvard University Press.

—— (2007). Structural Irrationality. In *Common Minds: Themes from the Philosophy of Philip Pettit*. Oxford University Press.

—— (2014). *Being Realistic about Reasons*. Oxford University Press.

Schroeder, Mark (2004). The Scope of Instrumental Reason. *Philosophical Perspectives*, *18*(1), 337–364.

—— (2007). *Slaves of the passions*. Oxford University Press.

—— (2008). Having Reasons. *Philosophical Studies*, *139*(1), 57–71.

—— (2009). Means-end Coherence, Stringency, and Subjective Reasons. *Philosophical Studies*, *143*, 223–248.

—— (2011a). What Does It Take to 'Have' a Reason? In *Reasons for Belief*. Cambridge University Press.

—— (2011b). Ought, Agents, and Actions. *Philosophical Review*, *120*, 1–41.

—— (2013). Scope for Rational Autonomy. *Philosophical Issues*, *23*(1), 297–310.

—— (2014). Hypothetical Imperatives: Scope and Jurisdiction.

—— (2015). Book Review: Being Realistic About Reasons, by Scanlon, T.M. *Australasian Journal of Philosophy 93(1)*: 195-198.

Setiya, Kieran (2007). Cognitivism about Instrumental Reason. *Ethics*, *117*(4), 649–673.

Shpall, Sam (2013). Wide and Narrow Scope. *Philosophical Studies 163,* no. 3: 717–36.

—— (2014). Moral and Rational Commitment. *Philosophy and Phenomenological Research 88,* no. 1: 146–72.

—— (forthcoming). The Calendar Paradox. *Philosophical Studies*.

Silk, Alex (2014). Why 'Ought' Detaches. *Philosopher's Imprint*.

—— (ms). *Discourse Contextualism: A Framework for Contextualist Semantics and Pragmatics*.

Skorupski, John (2010). *The Domain of Reasons*. Oxford University Press.

Strawson, Peter F. (1959). *Individuals: An Essay in Descriptive Metaphysics*. Routledge.

Strevens, Michael (2012). Ceteris Paribus Hedges- Causal Voodoo That Works. *Journal of Philosophy, CIX*(11).

—— (2013). Causality Reunified. *Erkenntnis*: 1–22.

Swanson, Eric. (2010). Lessons From The Context Sensitivity of Causal Talk. *The Journal of Philosophy*, *107*(5), 221–242.

—— (2010). On Scope Relations between Quantifiers and Epistemic Modals. *Journal of Semantics* 27 (4): 529–40.

—— (2012). The Language of Causation. In *The Routledge Companion to the Philosophy of Language*. Routledge.

Titelbaum, Michael (forthcoming). Rationality's Fixed Point (or: In Defense of Right Reason). In *Oxford Studies in Epistemology, Vol. 5*.

Way, Jonathan (2010). Defending the Wide-Scope Approach to Instrumental Reason. *Philosophical Studies*, *147*(2), 213–233.

—— (2010). The Normativity of Rationality. *Philosophy Compass*, *5*(12), 1057–1068.

—— (2011). Explaining the Instrumental Principle. *Australasian Journal of Philosophy*, 1–20.

—— (2013). Intentions, Akrasia, and Mere Permissibility. *Organon F*, *20*(4).

—— (forthcoming). Reasons and Rationality. In *Oxford Handbook of Reasons and Normativity*.

White, Roger (2014). Evidence Cannot Be Permissive. In *Contemporary Debates in Epistemology* (2nd ed.). Blackwell.

Williamson, Timothy (2011). 'Improbable knowing', in *Evidentialism and its Discontents* (147-164). Oxford University Press.

Worsnip, Alex (ms). The irrelevance of evidence (to rationality).

Wright, Crispin (2004). Warrant for Nothing (and Foundations for Free)? *Aristotelian Society Supplementary, 78,* 167–212.

Woodward, J. (2014) Scientific Explanation. *The Stanford Encyclopedia of Philosophy,* <http://plato.stanford.edu/archives/win2014/entries/scientific-explanation/>.