

Why Be Rational?

NIKO KOLODNY

Normativity involves two kinds of relation. On the one hand, there is the relation of being a reason for. This is a relation between a fact and an attitude. On the other hand, there are relations specified by requirements of rationality. These are relations among a person's attitudes, viewed in abstraction from the reasons for them. I ask how the normativity of rationality—the sense in which we 'ought' to comply with requirements of rationality—is related to the normativity of reasons—the sense in which we 'ought' to have the attitudes what we have conclusive reason to have. The normativity of rationality is not straightforwardly that of reasons, I argue; there are no reasons to comply with rational requirements in general. First, this would lead to 'bootstrapping', because, contrary to the claims of John Broome, not all rational requirements have 'wide scope'. Second, it is unclear what such reasons to be rational might be. Finally, we typically do not, and in many cases could not, treat rational requirements as reasons. Instead, I suggest, rationality is only apparently normative, and the normativity that it appears to have is that of reasons. According to this 'Transparency Account', rational requirements govern our responses to our beliefs about reasons. The normative 'pressure' that we feel, when rational requirements apply to us, derives from these beliefs: from the reasons that, as it seems to us, we have.

What are we saying when we say that, because of this or that, you 'ought' to have some attitude: to believe, desire, intend, or feel something? I think we are saying one of two things. Sometimes we are saying that there is *reason* for you to have the attitude, because of some feature of your *actual situation*. 'You ought to intend to leave; you have reason to! The building's on fire!' Other times we are saying that it would be *irrational* of you not to have the attitude, because of some *other attitude* that you have. 'Fine. Since that's what you intend, you ought to intend to open a new carton of cigarettes; it would be irrational of you not to. Whether or not there is any *reason* for you to smoke another one, you *intend* to smoke another one, and you know that you can't smoke another one unless you open a new carton'. When we say 'you ought to' in the sense of 'you have reason to,' we usually seem to be saying something about the relation between your *situation* and your *attitudes*. When we say 'you ought to' in the sense of 'it would be irrational of you not to,' we seem to be saying something about the relation between your *attitudes*, viewed in abstraction from the reasons for them. We are

saying something, for example, about whether your beliefs are logically consistent, or whether your intentions for ends cohere with your intentions for means—things that are true, if they are, quite independently of whether there is reason for you to have any of those beliefs or intentions. Now, I don't claim that everyone uses the words 'reasons' and 'rationality' to distinguish these two kinds of normative relations. In particular, there is a common use of 'rational'—we might call it the 'objective' use—in which the phrase 'the rational thing for one to do' simply means, in my vocabulary, what one has most reason to do. To avoid this confusion, it might be better to call normative relations of the second kind relations of '*subjective*' rationality.¹ But whatever words we use to describe it, my point is there is an important difference between these two normative relations. My topic is what the relation between these normative relations, of reasons and of rationality, is.

Some—call them 'rationalists'—aim to explain reasons in terms of rationality. They 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. And they then understand reasons for action in terms of what such an agent would desire or will. Smith (1994), for example, proposes that one's having a reason to do something in given circumstances is its being the case that one would desire that one did it in those circumstances, if one were fully rational.² It is not immediately obvious how Smith, having analysed the 'ought' of reasons in terms of the 'ought' of rationality, then proposes to understand the 'ought' of rationality. But one imagines that Smith will deny

¹ In the case of belief, it is crucial to distinguish objective epistemic rationality from subjective epistemic rationality, on the one hand, and from true belief, on the other. Objective epistemic rationality consists, roughly, in having the beliefs that the evidence in fact supports. These beliefs may be false; the evidence may mislead. So objective epistemic rationality differs from true belief. There are also certain requirements of theoretical rationality that one can satisfy without having beliefs that the evidence supports. One might believe that *p* and that *p* entails *q*, defying overwhelming evidence to the contrary. Still, in so far as one has these beliefs, it is irrational of one not to believe that *q*. In forming the belief that *q*, one avoids this kind of irrationality. But that does not mean that the evidence supports that *q*. So subjective epistemic rationality differs from objective epistemic irrationality. I am not proposing, here, a substantive account of what subjective epistemic rationality is. I am just suggesting that it is plausible that there are some rational requirements that one can satisfy by having beliefs that one has no reason to have. On the substantive account I favour, subjective epistemic rationality consists in having the beliefs that one *believes* that the evidence supports. See principles B+ and B–, below, and the principles from which I claim that they are derived, C+ and C–, at the end of this paper.

² For Smith, being fully rational involves, in part, being fully informed. And being fully informed is not simply a matter of structural or procedural relations among one's attitudes. It consists in there being a certain relation between the facts and one's beliefs. But, still, this relation is not the relation of being a reason for doing such and such, or even that of being a reason for believing that such and such. So, to the extent that Smith intends to offer an analysis of what is for someone to have a reason to do something, there is no circularity.

that there is any further question here. What more can it mean that we ought to *X*—he might say—but that it would be irrational of us not to *X*?

Kantians also seek to explain reasons in terms of rationality, at least in broad outline. Korsgaard (1996) traces the ‘source’ of normativity to rational agency. By willing in conformity with the principles of rationality, we confer the status of being a reason on features of the world. This is a somewhat misleading formulation, of course, since, for Korsgaard, there is a sense in which willing in conformity with the principles of rationality just is willing itself. The principles of rationality are constitutive of the will; it is only by conforming to them that one counts as having a will at all. And this, she argues, explains why there can be no real question ‘Why ought I to conform to principles of rationality?’ Not to conform to them just is, at the limit, to cease to have a will—to cease to be someone for whom the question, ‘Why ought I?’ can even arise.

By contrast, ‘nonreductionists,’ such as Broome (2004), Dancy (2000), Parfit (ms), Raz (1999), and Scanlon (1998), deny that reasons for action, and perhaps also reasons for belief, can be explained in terms of rationality, or indeed in any other terms.³ The relation of being a reason for is primitive. This view provokes several concerns, which many think that rationalism is better placed to address. Some of these concerns are metaphysical. What is it for a fact to be a reason for an attitude? Taken at face value, it seems a murky relation between world and mind. Rationality, by contrast, is simply a matter of having certain psychological patterns, and we know what those are. Other concerns are epistemological. How can we settle questions about what is a reason for what? As far as the bare concept of a reason is concerned, it seems that anything might be a reason for anything else. By contrast, it seems uncontroversial that rationality is a kind of coherence or unity. So it is relatively clear how we might settle questions about what rationality requires; it is whatever is necessary for coherence. Still other concerns have to do with motivation. Why is it the case that the judgement that I have reason to *X* will, in so far as I am rational, motivate me to *X*? If that judgement is treated as an unanalysed primitive, then it is left mysterious. But if we explicate that judgement in terms of the concept of rational motivation itself, then we appear to be on our way to an explanation.

My topic in this paper is not these familiar concerns, although I believe that the account of rationality that I will sketch helps to address

³ Or at least that the ‘ought’ of reasons can be explained in other terms.

them. My topic is instead a different concern about nonreductionism, which has not been much discussed. If nonreductionism takes the 'ought' of reasons to be primitive, then how does it understand the 'ought' of rationality? Nonreductionists do not deny that there are requirements of subjective rationality: requirements to do what one believes one ought, or to take the means to one's ends. Nor can they deny that we ought, in some sense, to comply with these requirements. But how is this 'ought' of rationality to be understood? Must nonreductionists take it to be a second primitive? The worry about this option is not simply the proliferation of primitives.⁴ It is also the difficulty of understanding how one and the same subject is to be governed by these two autonomous 'ought's, which sometimes issue incompatible directives. Alternatively, nonreductionists might try to explain the 'ought' of rationality in terms of the 'ought' of reasons, just as rationalists propose to explain the 'ought' of reasons in terms of the 'ought' of rationality. The question is whether nonreductionists can do this.

To illustrate just one of several difficulties, suppose I believe that I have conclusive reason to have some attitude. In some sense, I ought to have that attitude; it would be irrational of me not to have it. Now suppose that 'ought' here means 'have reason.' Then we get the bootstrapping result that if I believe that I have conclusive reason to have some attitude, then I in fact have reason to have it. This is absurd. As I argue in section 1, an otherwise attractive attempt to overcome this bootstrapping problem is unsuccessful. Rational requirements cannot be understood as 'wide-scope' requirements, as John Broome suggests, because not all rational requirements are wide scope.

Moreover, two further arguments, which are independent of the bootstrapping problem, bolster the conclusion that we do not have reasons to be rational. As I argue in section 2, it is unclear what reasons there might be for complying with rational requirements. And, as I argue in section 3, even if we could identify some putative candidates, they typically would not, and in some cases could not, function as reasons for us. We typically would not, in some cases could not, reason from them.

⁴ It might seem odd, from a certain vantage, for nonreductionists to care about such parsimony. If they accept one mysterious primitive, one might wonder, then why should they be reluctant to accept another? However, nonreductionists' willingness to accept reasons as primitive follows a long history of failed attempts to reduce the 'ought' of reasons to something nonnormative. There is no comparable history of attempts to reduce the 'ought' of rationality to the 'ought' of reasons. Moreover, there are *prima facie* grounds for greater optimism about this latter kind of reduction. A reduction of one 'ought' to another 'ought,' of one part of the normative to another part, seems more likely to succeed than a reduction of an 'ought' to something nonnormative.

The upshot of these three arguments is that the normativity of rationality is not straightforwardly the normativity of reasons. All the same, as I argue in section 4, we cannot rest content with the idea that rationality is purely evaluative and deny that it is in any sense normative. The requirements of rationality may be, in part, standards of appraisal, by which we measure how far someone manifests a kind of virtue or proper functioning. But they also seem to call for certain responses. When it would be irrational of us to fail to have some attitude, we typically feel that we ought to have it. Furthermore, claims about rationality are sometimes put forward as a kind of advice. But how can this be, if we don't have reason to comply with rational requirements?

What I suggest, in section 5, is that rationality is only apparently normative, but in a way that gives a function to this special kind of advice that employs the 'ought' of rationality. Suppose—and I grant that this is no small supposition—that all rational requirements take a certain form: namely, that all rational requirements require one either to have the attitudes that one believes that one has conclusive reason to have, or to lack the attitudes that one believes one lacks sufficient reason to have. This supposition makes possible the following 'Transparency Account.' When we advise someone that he 'ought rationally' to have some attitude, we aim not to offer him a reason that *we* believe has for that attitude—as we do with normal advice—but instead to draw his attention to a reason that *he* believes he has. We are saying, in effect: 'As it seems to you, you have reason to have that attitude.' The normative pressure, so to speak, that the advisee then feels to comply with the rational requirement, by forming the attitude, derives from how things seem to him—from the reason that, as it appears to him, he has. Thus, while the fact that a rational requirement applies to one is not *in fact* a reason for one to comply with it, it will always *seem* to one, when one is subject to a rational requirement, that one has a reason of another kind to comply with it: namely, the reason to form (or drop) the attitude that, in so far as one satisfies the antecedent of the requirement, one already believes one has (or lacks). This is what gives the 'ought' of rationality its normative force—or, rather, its seeming normative force.

Notice that this account does not appeal to any second primitive concept, beyond that of a reason. So while rationality is not normative in the sense that there are reasons to be rational, the sense in which rationality is normative—or, better, apparently normative—is explained in terms of the concept of a reason. In this sense, then, the Transparency Account explains rationality in terms of reasons.

1. First argument that there are no reasons to be rational: bootstrapping

Consider the view of the normativity of rationality once held by Broome: that there are conclusive reasons to comply with rational requirements.⁵ Why doesn't this lead to bootstrapping? Because, Broome argues, rational requirements have 'wide scope.' Contrast two forms that rational requirements might take. According to the first, rational requirements have narrow scope: that is, when one has attitude *A*, then one is rationally required to have attitude *B*. According to the second, rational requirements have wide scope: that is, one is rationally required (if one has attitude *A*, to have attitude *B*). Now consider the loosely stated rational requirement: 'If you believe that *p*, then you are rationally required to believe what *p* entails.' If we read this literally, and assume that the rational requirement has narrow scope, then 'detachment' of the consequent is permitted. From the fact that you believe that *p*, it follows that you are rationally required to believe what *p* entails. If we suppose that we have conclusive reasons to comply with rational requirements, then it follows that you have conclusive reason to believe what *p* entails. And since *p* entails *p*, it follows that if you believe that *p*, then you have conclusive reason to believe that *p*. Beliefs become, as Broome (1999, p. 405) puts it, 'self-justifying.'⁶ We avoid this result, Broome observes, if we interpret the rational requirement as

⁵ Broome is presently revising his view of the normativity of rationality, so nothing here should be taken to represent his current thinking. The view I discuss can be found in Broome (1999), (2001), (2003a), (2003b), and (2004). Broome (1999), for example, understands rational requirements as a species of 'normative requirements,' and uses the same 'ought' to characterize both normative requirements and conclusive reasons for particular attitudes. The 'logical factor' of the proposition that *p* is a conclusive reason for one to make it the case that *q*, Broome claims, is if *p*, then *Oq*. By contrast, the logical factor of the proposition that *p* 'normatively requires' one to make it the case that *q* is *O*(if *p*, then *q*). The 'O' or 'ought' operator is the same. The only difference is whether it governs the consequent, or the whole conditional. Broome (2004, p. 51) makes the point explicitly: 'The fact that *q* follows obviously from *p* explains why you ought (to believe *q* if you believe *p*). So this fact constitutes a perfect reason for you (to believe *q* if you believe *p*).' Broome (ms) abandons this view, on grounds similar to those I raise in section 2. Dancy (2000, pp. 60–70) proposes an account similar to Broome's original view. Yet Dancy suggests, rightly in my view, that this account must be 'supplemented' with the idea that subjective requirements are not only 'deontic,' but also 'evaluative' (p. 64). Claims about rationality are sometimes—although, significantly, not always—in the register of assessment, rather than that of advice. I discuss this in greater detail in section 4.

⁶ To say that beliefs become 'self-justifying' may be misleading. The argument does not show that there is conclusive *epistemic* reason for every belief. For that conclusion, we need the further premiss that the 'ought' in question is supported by epistemic reasons. As I argue below, reasons to conform to rational requirements may have to be understood not as epistemic reasons, but instead as reasons to avoid irrationality as a kind of intrinsic bad. Nevertheless, the conclusion that there is conclusive reason, of some kind or other, for every held belief seems sufficiently absurd for Broome's purposes.

having wide scope: that is, as ‘You are rationally required (if you believe that p , to believe what p entails).’ If rational requirements have this form, then detachment is not permitted. From the fact that you believe that p , it does not follow that you are rationally required to believe what p entails. So even if you have conclusive reason to have the attitudes that you are rationally required to have, it does not follow that you have conclusive reason to believe what p entails. In sum, understanding rational requirements as having wide scope leaves us free to understand the normativity of rationality in terms of reasons. We can have conclusive reasons to comply with rational requirements without bootstrapping.

The problem with Broome’s approach is that some rational requirements have narrow, rather than wide, scope. If we understand the normativity of rationality in terms of reasons to comply with rational requirements, therefore, then we license bootstrapping. It would follow, for example, that merely believing that I had conclusive reason to X would suffice to give me reason to intend to X .

1.1 What rational requirements are

Before arguing for this, however, I need to clarify what rational requirements are. I agree with Broome that it is often convenient to describe features of subjective rationality in terms of rational requirements. But we need to treat this talk of rational requirements with some caution. In real life, we don’t ever say that ‘rationality requires’ you to X . It is a philosopher’s phrase of art. One way or another, we must anchor this talk in terms of the more familiar judgements that we make about people’s rationality and irrationality. Roughly, ‘ S is rationally required to X ,’ ought to be equivalent to more familiar judgements such as, ‘ S is, or would be, being irrational in not X -ing,’ or ‘It is, or would be, rational for S to X .’

Our ordinary judgements about rationality have two important features, which rational requirements should reflect. First, these judge-

Yet some might disagree. Harman (1986) suggests that the mere fact that one has a belief, at least in the absence of any contrary evidence, provides some reason, indeed some epistemic reason, for one to have it. I readily grant that believing that p , when there is evidence that one is reliable in the domain in question, can be evidence that p . But *merely* believing that p , in the absence of any evidence that one is reliable in the domain in question, is not evidence that p . Compare Pollock and Cruz (1999, Ch. 3) and Jones (ms a). At any rate, even if there are some who embrace bootstrapping reasons for belief, I do not know of anyone who embraces bootstrapping reasons for action. It seems generally accepted that merely believing that one has conclusive reason to X does not give one reason to X .

ments are local. They are focused on specific conflicts among one's attitudes. We might judge, for example, that a person is being weak-willed in believing that he has conclusive reason to *X*, but not intending to *X*. And we might judge, at the same time, that he is giving in to wishful thinking in believing that he has conclusive reason to *X*, having himself decided that the evidence for that belief is flimsy.⁷

Rational requirements, accordingly, ought to be local. In each instance in which one is under a rational requirement, what it ought to require of one is to avoid or resolve some specific conflict among one's attitudes—as opposed to, say, to satisfy some global constraint on all of one's attitudes. One might liken each application of a rational requirement to a referee with authority over a different part of a playing field, or to an inspector with authority over a different stage in a production process. Various applications of rational requirements may call for one to form some attitudes, to retain others, and to revise still others all at the same time.⁸

The other important feature of our ordinary attributions of rationality and irrationality is that they attach not only to states, but also to

⁷ Judgements of someone's being rational or irrational, like judgements of someone's having or lacking reason, are fundamentally relational. They bind an agent and a set of attitudes, or an agent and a transition between sets of attitudes. Hence, we say that someone is being irrational *in* intending the end, but not the means, or that it is irrational for him *to* intend the end, but not the means. One might think that these two-place predicates—'S is irrational in X-ing' and 'it is irrational for S to X'—ought to be analysed in terms of the one-place predicate 'S is irrational.' But I think that this is a mistake. One wouldn't attempt to analyse the two-place predicate, 'has reason to,' in terms of the one-place predicate, 'has reason.' Indeed, I think that the one-place predicate is simply a degenerate instance of the two-place predicate. 'Is irrational' means 'is irrational in having some attitudes or other.' This is so at least if 'is irrational' is predicated of a person at a particular time. 'Is irrational' can also be used to ascribe a long-term disposition. In this case, it means that the person is disposed to be irrational in the first sense, which in turn means that he is disposed to be irrational in having certain attitudes or other.

⁸ Broome (ms) agrees. Others, however, seem to take a different view. Although they start with local principles of rationality, they go on to suggest that local requirements of rationality are only *prima facie* requirements. What rationality requires of one 'all things considered,' as it were, is some function of all of these local requirements: presumably, what will maximize one's overall satisfaction of these various local requirements, perhaps with weightings for more and less important local requirements. I am sceptical of this idea. For one thing, it does not correspond to any recognizable feature of our ordinary attributions of subjective rationality and irrationality. Moreover, it is obscure what role such 'all things considered' requirements could play. Unless I am under conflicting local requirements, I maximize compliance with local requirements simply by complying with each of them, in which case a requirement to maximize compliance is superfluous. A requirement to maximize compliance overall would have to be a principle of triage, determining which from among conflicting requirements I ought to satisfy. It isn't clear to me that there is, or needs to be, a determinate answer to the question of which from among conflicting requirements one ought to satisfy.

processes. We judge that a person is rational or irrational not only in virtue of the state he is in at a given time, but also in virtue of how he transitions from one state to another over time. A bolt of lightning might jolt me out of a state in which I have two inconsistent beliefs and into a state in which I lack one or both of them. Although this process might be said to take me from an irrational state to a rational one, it would not, itself, be rational. By contrast, it would be rational of me to revise one or both of these beliefs in light of the evidence against them. In other words, one is rational or irrational not only in virtue of the attitudes that one has at any given moment, but also in virtue of how one forms, retains, and revises one's attitudes over time.

We should, accordingly, distinguish between 'state-requirements,' which simply ban states in which one has conflicting attitudes, and 'process-requirements,' which say how, going forward, one is to form, retain, or revise one's attitudes so as to avoid or escape such conflict-states. Any account of rationality that aims to capture our ordinary attributions cannot consist solely of state-requirements, which say 'Not that conflict!' It must consist, at least in part, of process-requirements, which say 'Do this to avoid or resolve that conflict!'

My argument that some rational requirements have narrow scope, as we will see, needs only the weaker claim that some rational requirements are process-requirements. I am inclined to think, however, that all rational requirements are process-requirements. Being rational just is responding in the ways that process-requirements call for. To begin with, notice that once we have specified the process-requirements of rationality, state-requirements become superfluous. By complying with process-requirements, one thereby complies with any state-requirements that might apply to one. If one *avoids* or *escapes* conflict-states, then one won't be *in* them. Moreover, the very idea of a state-requirement is questionable. If rational requirements are normative, deontic, or response guiding, then they call for the subject to *respond* in a certain way. It is clear how forming, retaining, or revising one's attitudes so as to *avoid or escape* a conflict-state might qualify as a response. But how might *not being in* a conflict-state qualify as a response? Indeed, one feels driven to interpret the claim that one is required not to be in a conflict-state as simply the claim that one is under a very general process-requirement: to avoid or escape that conflict-state in any way one likes. State-requirements do not seem to be normative; they do not say what we are to do. At most, they seem evaluative; they seem to be merely standards for appraising our current condition. In so far as gen-

uinely normative requirements of rationality are concerned, the only such requirements appear to be process-requirements.⁹

Broome seems to take a different view: that all rational requirements are state-requirements. However, this difference may be less significant than it at first appears. For Broome acknowledges that we need to explain how we can rationally bring ourselves into conformity with state-requirements, and his pioneering work on reasoning is meant to answer this need. In any event, the crucial point is that Broome needs some account of how we can rationally bring ourselves into conformity with state-requirements. For rationality is, in no small part, a matter of how one forms, retains, and revises one's attitudes.

1.2 Motivating the 'Reasoning Test' for the scope of rational requirements

So rational requirements are local, and some, or all, are process-requirements. Now we come to Broome's claim that rational requirements have wide, rather than narrow, scope. Why should we accept it?

One possible argument that rational requirements have wide scope is the following *reductio*. If rational requirements have narrow scope, and if we have conclusive reasons to comply with rational requirements, then there is bootstrapping. But, clearly, there cannot be bootstrapping. And we have conclusive reasons to comply with rational requirements. Therefore, rational requirements do not have narrow scope. The problem with this argument is that it assumes what is ultimately in question: namely, that we have conclusive reasons to comply with rational requirements.¹⁰

A different, but equally unsatisfactory argument is that narrow-scope requirements are ruled out by the very nature of subjective rationality. Subjective rationality is a matter of the relations among one's attitudes.

⁹ I do not mean to deny that rational requirements are organized around specific kinds of conflict-states. The basic schema for a rational requirement is: Do *X* to avoid or resolve conflict-state *Y*. Accordingly, we typically identify rational requirements by first identifying conflict-states to be avoided or resolved. This is, in essence, what it means to say that rational requirements are 'local.' Nor do I mean to deny, at least not here, that there may be evaluative standards according to which it is, in some sense, bad to be in a conflict-state. It might be argued that these standards give the normative requirements of rationality their point. I am grateful to Luca Ferrero and Pam Hieronymi for urging me to clarify this.

¹⁰ One might try to make do without the claim that we have reasons to comply with rational requirements. Compare Broome (2003b, lec. 1, p. 2). Take a narrow-scope requirement such as: if I believe that I have conclusive reason to *X*, then I am rationally required to intend to *X*. Suppose that I believe that I have conclusive reason to *X*. Then it would follow that I am rationally required to intend to *X*. But that is absurd. *X*-ing, and intending to *X*, might be completely silly, or worse. If one does find it absurd, however, then one is either tacitly assuming the claim that we have conclusive reasons to comply with rational requirements, or one is reading 'rationally required' as 'objectively rationally required.' In either case, it follows that I have conclusive reason to intend to *X*. If

Therefore, rational requirements govern the relations among one's attitudes. Therefore, rational requirements have wide scope. This last step, however, is a *non sequitur*. A narrow-scope requirement, which says that if one has attitude *A*, then one is rationally required not to have attitude *B*, also governs relations among one's attitudes.¹¹

A more promising argument starts from the observation that state-requirements ban patterns of attitudes, such as having both *A* and *B*. More than one state will satisfy such a ban. The state of not having *A* satisfies it, as well as the state of not having *B*. So the state of not having *A* is not rationally required. What is rationally required is (either not having *A*, or not having *B*). In other words, this state-requirement has wide scope.

The question is whether this argument about *state*-requirements can be extended to an argument about *process*-requirements. At first glance, the extension might seem straightforward. If two states provide alternatives to a conflict-state, one might argue, then there are two ways of revising one's attitudes so as to escape that conflict-state. One can revise one's attitudes so as to enter into the first alternative state, or one can revise one's attitudes so as to enter into the second alternative state. Suppose that one has conflicting attitudes *A* and *B*. One is rationally required to resolve this conflict. There are two ways of doing this. One can revise *A*, or one can revise *B*. So one is not rationally required to resolve the conflict by revising *B*. Instead, one is rationally required (either to revise *A*, or to revise *B*).¹²

This is the argument, it seems, on which the claim that process-requirements have wide scope must ultimately rest. The argument relies crucially on the premiss that there is more than one rational way to resolve the conflicts that rational requirements govern. This gives us a test of the scope of process-requirements. Suppose it is claimed that the process-requirement governing the conflict between *A* and *B* is

X-ing, and intending to *X*, are completely silly, then that implication is indeed absurd. But, again, one cannot assume, at this point, that we have conclusive reasons to comply with rational requirements. And 'rationally required,' in the present context, means 'subjectively rationally required,' not 'objectively rationally required.'

¹¹ Compare Broome (2003b, lec. 1, p. 4).

¹² Compare Broome (2001, p. 180): 'Reasoning is correct if it makes your mental states conform to normative requirements you are under. You can conform to a normative requirement such as [the one governing means-end rationality] in two ways. You may enter the conclusion-state or alternatively you may leave one of the premiss-states; you may give up one of your existing beliefs or intentions A process that has either of these results will be correct.'

wide scope: i.e., one is rationally required (either not to have *A*, or not to have *B*). For this claim to be true, it must be the case that (i) one can rationally resolve the conflict of having *A* and *B* by dropping *B* and (ii) one can rationally resolve it by dropping *A*. Call this the ‘Rational-Response Test.’

To refine this test, we need to clarify what ‘rationally resolving’ a conflict comes to. Suppose that dropping *B* is (at least) one way in which one can rationally resolve the conflict of having both *A* and *B*. Not every case of dropping *B*—not every case in which the state (*A*, not *B*) follows the state (*A*, *B*)—will be a case in which one actually does rationally resolve the conflict. For example, one might drop *B* as a result of an electric shock. Why isn’t this a rational resolution of the conflict? Because, roughly, one’s awareness of what is amiss in the state (*A*, *B*) does not explain the transition to (*A*, not *B*).

The question is how to understand this awareness. It cannot be that the subject reflects on his attitudes *themselves*, recognizes that his attitudes violate a rational requirement, and then makes the appropriate adjustments on that basis. As I will go on to argue, we typically do not comply, and in some cases cannot comply, with rational requirements in this way. How else, then, is the subject’s awareness to be understood? The alternative, I think, is this. From the standpoint of attitude *A*—which has at its object the *content* of *A*, not *attitude A* itself—the subject is aware of a need to revise his other attitude, *B*. Then, on the basis of the content of attitude *A*, the subject revises *B*. In a broad, but recognizable, sense of ‘reasoning,’ the subject *reasons* from the content of *A* to revising *B*. Suppose, for example, the conflict consists in my believing that it is Monday and believing that it is Tuesday. If I rationally resolve this conflict by revising my belief that it is Tuesday, I do so by reasoning from the content of my belief that it is Monday—from the fact, or apparent fact, that it is Monday—to revising the belief that it is Tuesday. I don’t mean to suggest that this reasoning is explicit, or deliberate. In the vast majority of cases, our reasoning—our acquiring, retaining, and revising attitudes on the basis of the contents of other attitudes—is implicit and not voluntarily directed.¹³ We can thus restate the ‘Rational-Response Test’ as the ‘Reasoning Test.’ The proc-

¹³ By contrast, Broome (2003b, lec. 1 pp. 8–9) reserves the title, ‘reasoning,’ for a special case: a deliberate, explicit process—indeed, a complex action—in which one expresses certain attitudes to oneself in order to cause oneself to acquire, retain, or lose certain other attitudes. He doubts that one can come to satisfy all rational requirements (which he takes to be state-requirements) by reasoning, so understood. Nevertheless, he seems willing to accept that even if one cannot satisfy all rational requirements by reasoning in this narrow sense, one can still satisfy them by ‘passive processes,’ which include ‘reasoning’ in the broader sense of acquiring, retaining, or losing certain attitudes on the basis of certain others.

ess-requirement governing the conflict between *A* and *B* is wide scope—that is, one is rationally required (either not to have *A*, or not to have *B*)—only if, from a state in which one has conflicting attitudes *A* and *B*, (i) one can reason from the content of *A* to dropping *B* and (ii) one can reason from the content of *B* to dropping *A*.

1.3 Rational requirements governing conflicts between attitudes and beliefs about reasons for them

Some rational requirements may pass this test. But do all? Not, I think, process requirements that govern conflicts between attitudes and beliefs about reasons for them, such as:

- B+: Rationality requires one to believe that *p*, if one believes that there is conclusive evidence that *p*.
- I+: Rationality requires one to intend to *X*, if one believes that there is conclusive reason to *X*.
- B–: Rationality requires one not to believe that *p*, if one believes that there is not sufficient evidence that *p*.
- I–: Rationality requires one not to intend to *X*, if one believes that one lacks sufficient reason to *X*.¹⁴

In the following two subsections, I will argue that the Reasoning Test shows that these requirements have narrow scope. If the reader finds these requirements plausible, then I encourage him or her to skip ahead. However, some readers may have concerns about these requirements. In the rest of this subsection, I try to address these concerns.

To begin with, one might doubt that *every* such belief about reasons for an attitude triggers a rational requirement. Suppose a blow to the head gives me the belief that there is conclusive evidence that *p*. Is it really true that I am rationally required to believe that *p*, or—not to prejudice the question whether the requirement would have narrow scope—rationally required (either not to believe that there is conclusive evidence that *p*, or not to believe that *p*)? It would not undermine the argument of this section to concede that beliefs about reasons caused in this way do not trigger rational requirements. So long as one accepts that beliefs with the same content, but arrived at in other ways, trigger rational requirements, and that these beliefs are sometimes false,

¹⁴ Broome (2001) explicitly recognizes I+ as a rational requirement. Broome (2003b, lec. 2, p. 4) also appears to endorse requirements like B+ and B–. For the most part, however, his examples of requirements of theoretical rationality do not involve beliefs about the evidence as such.

then one can agree with the conclusion of this section: that there are not conclusive reasons to comply with rational requirements in general, because, in at least some cases, this would lead to bootstrapping. As it happens, restricting the set of relevant beliefs in this way would be incompatible with the Transparency Account of section 5. Yet the Transparency Account helps to alleviate the doubt that motivates the restriction in the first place, for it explains why the normative ‘pressure’ that a subject feels to comply with the ‘ought’ of rationality should be present, and the kind of advice couched in terms of that ‘ought’ that others can give him should be available, whenever a subject has a belief that he has conclusive reason, even a belief resulting from brute force. A sufficient response to the present doubt, therefore, may be to say, first, that it doesn’t matter for the purposes of this section whether we let the doubt stand and, second, that later discussion may help to remove it, by situating it within a broader view of the nature of rational requirements.

In any event, we may be able to allay this doubt by using the resources already at our disposal. The doubt may partly stem from the thought, with which I agree, that a belief produced by a blow to the head does not entail anything about what further attitudes one has reason to have. But B+ would entail this only with the further assumption that there are reasons to comply with rational requirements, which is precisely what is at issue. We cannot determine which rational requirements there are, in the present context, by assuming that we have reasons to comply with them. Instead, we must consider, as discussions of rationality commonly do, when there is a conflict among attitudes that attracts attributions of irrationality, which are ordinarily expressed by such formulations as, ‘S is being irrational in not X-ing.’ It seems hard to deny that there is at least one such conflict in the mind of someone who believes that there is conclusive evidence that *p*, but does not believe that *p*, even when the former belief is the product of brute force. It might be odd to focus on this conflict in such a case, given that something else is amiss that is far more salient: namely, that one of the conflicting beliefs resulted from a blow to the head. But this does not mean that the conflict is not there, or that by letting it stand the subject is not being in at least one way irrational, whatever other, more striking aberrations in his flow of thought there may be.

As one reader has observed, it would be even harder to deny that B+, B–, I+, and I– were valid requirements of some kind if they were conditioned not on *beliefs* about reasons, but instead on the *reasons* themselves. If there is conclusive evidence that *p*, then one is, in some sense,

‘required’ to believe that p . The sense in which one is ‘required’ to believe it, I think, just is that there is conclusive evidence that it is so, that one has conclusive reason to believe it. This ‘objective’ requirement, in other words, simply restates its condition. Now, such objective requirements are not themselves requirements of subjective rationality, since they do not govern relations among attitudes, in abstraction from the reasons for them. But there does seem to be an intimate connection between these objective requirements and the four requirements of subjective rationality that we have been considering. The latter, one might say, are the psychological shadows cast by the former; the subjective requirements govern the progress of thought of someone who believes that she is bound by the objective requirements. And this may prompt the suggestion that the objective requirements might somehow be basic, with these four subjective requirements being extensions of a kind from them. In fact, the Transparency Account represents one way of cashing out this suggestion. But it would be premature to take up this suggestion now.

A different worry arises from a set of potential counterexamples to B+, B−, I+, and I−. One might believe that there is conclusive evidence that p , for example, but believe that it would be drastically worse to believe it. Would it necessarily be irrational of one not to believe that p , as B+ claims? Or one might believe that there is not sufficient evidence that p , but believe that it would be overwhelmingly better to believe it. Would it necessarily be irrational to believe it, as B− claims? Similarly, one might believe that one had conclusive reason to *do* X , but not believe that one had conclusive reason to *intend* to do X . For example, one might believe that intending to X is unnecessary for X -ing,¹⁵ or prevents one from X -ing.¹⁶ Would it be irrational of one not to intend to do X , as I+ claims? Or one might believe that one lacked sufficient reason to do X , but not believe that one lacked sufficient reason to intend to do X , because one believed that one had reason to intend to X that was independent of one’s reason to X . An eccentric might offer one a million dollars, for example, to intend to drink a toxin here and now,

¹⁵ As Kamm (2000) argues, one may do X , without intending it, as by-product of doing some Y that one intends. To alter slightly an example that Broome (2003b, lec. 1) uses to illustrate Kamm’s point, one might believe that one has conclusive reason to have more fun and that one has conclusive reason to cultivate new friends. In order to have more fun, one might intend to buy a boat. One might realize that by doing so, one will cultivate one new friends, whether or not one intends to. In this case, Broome suggests, it does not seem irrational of one not to intend to cultivate new friends. In an even wider range of cases, one may believe that forming now—that is, upon judging that one has conclusive reason to X some time in the future—an intention to X is not necessary for X -ing. It may lie so far in the future that one doesn’t need to decide this very moment. In this case, it does not seem irrational of one not to intend, now, to X

whether or not one goes on to drink it.¹⁷ Would it be irrational, as I— claims, to intend to *X*, supposing that it was in one's power to intend it?

If one finds the counterexamples to B+ and B— compelling, I suspect, it is because one thinks that there can be, and that subjects can in practice take there to be, reasons for believing that *p* that are not evidence that *p*. And if one finds the counterexamples to I+ and I— compelling, it is because one believes that there can be, and that subjects can in practice take there to be, reasons for intending to do *X* that are not reasons for doing *X*. One's underlying complaint, in other words, is that B+, B—, I+, and I— restrict the class of relevant reasons for belief and intention unnecessarily. This complaint could not be raised against the following requirements, which simply leave open what can count as a reason for belief or intention:

C+: Rationality requires one to have *A*, if one believes that one has conclusive reason to have *A*; and

C—: Rationality requires one not to have *A* if one believes that one lacks sufficient reason to have *A*,

with 'to have *A*' replaced by 'to believe that *p*,' or 'to intend to do *X*.' My reply to someone who finds the counterexamples compelling, therefore, is simple: substitute C+ and C— for B+, B—, I+, and I— in what follows. I have opted to present the argument in terms of B+, B—, I+, and I—, since I expect that most readers will find them more natural. But the argument of the present section does not require it.

Although it does not affect the present argument, I am not agnostic on these matters. As we will see, the Transparency Account requires that C+ and C—, or the 'core requirements,' are in fact the basic requirements of rationality. All other rational requirements must either be derived from the core requirements, or be valid, when they are, only by approximating them. The counterexamples to I+ and I—, I believe, help to show that I+ and I— are valid, when they are, only when they approximate C+ and C—. Typically, if one has reason to *X*, then the fact

¹⁶ Intending to be more spontaneous, Jonathan Dancy notes, may prevent one from being spontaneous. Similarly, Broome (2001) observes, intending to sleep may prevent one from sleeping (although, in this case, Broome believes that it *is* irrational for one not to intend to sleep). One might object, not unreasonably, that one cannot intend to be spontaneous or to fall asleep in the first place, because these are not actions: things that one can intentionally do. But similar examples seem less vulnerable to this objection. Suppose that being spontaneous amounts to doing things on whims, without making plans in advance. While doing something on a whim is something that one can do intentionally, it would be self-defeating to form a prior intention to be spontaneous in this sense.

¹⁷ A variant of the famous example of Kavka (1983).

that intending to *X* is a constituent of, or necessary means to, *X*-ing provides reason of the same stringency to intend to *X*. And, typically, there is no other reason to intend to *X*. So, typically, when one believes that one has conclusive (or lacks sufficient) reason to *X*, one also believes that one has conclusive (or lacks sufficient) reason to intend to *X*. So, typically, when one satisfies the antecedent of *I*+ (or *I*-), one also satisfies the antecedent of *C*+ (or *C*-), where 'to have attitude *A*' is replaced by 'to intend to *X*'. But, as the counterexamples show, there are atypical cases. If one believes that one is in such an atypical case, then one can satisfy the antecedent of *I*+ (or *I*-) without satisfying the antecedent of *C*+ (or *C*-). When this occurs, it does not seem irrational of one not to intend to *X* (or to intend it). This suggests that the genuine requirements of rationality are *C*+ and *C*-, with *I*+ and *I*- being, as it were, rules of thumb. By contrast, I am not convinced by the counterexamples to *B*+ and *B*-. While there are no compelling grounds for thinking that reasons for intention just are reasons for action, or, more to the point, that subjects must in practice treat them as identical,¹⁸ there are compelling grounds for thinking that reasons for belief just

¹⁸ Some may think that the toxin puzzle of Kavka (1983) provides such grounds, but it is not clear to me how it does. To begin with, nothing in the puzzle touches the claim that reasons for intention are distinct from, although always instrumentally related to, reasons for action. This claim would suffice for my purposes. Now, one might grant this claim, but argue that the puzzle still shows that all reasons for intention are instrumentally related to reasons for action: that the only reason for intending something is that it is instrumental to doing something that one has reason to do. But it does not show even this. It shows only that, in certain situations, one cannot intend to do what one believes that one has no reason to do. Even if it is true that one cannot have reason to intend what one cannot intend, it would only follow that in these situations one cannot have reason to intend. Nothing would follow about other situations. At the heart of the toxin puzzle is the following constitutive feature of intention: that one intends to *X* only if one believes that it is possible that one will *X*. In the puzzle, one is asked to intend at midnight to drink the toxin tomorrow afternoon, and if one so intends, one will be rewarded hours before the time comes to drink. Kavka observes, quite plausibly, that it is doubtful that one could intend to drink. The reason is this. One foresees that, after midnight, but before the time comes to drink, one will recognize that one has no reason to intend to drink and, indeed, compelling reason not to intend. One also foresees that one will be sufficiently rational tomorrow to revise one's intentions accordingly. Therefore, one foresees that no intention one forms at midnight can survive until the time comes to drink. Since one believes that one will drink only if one intends to, one believes that one simply will not drink. Since one cannot intend what one believes it is not possible for one to do—such is the constitutive feature of intention—one cannot intend to drink. This problem, however, does not present any categorical barrier to intending to do what one believes one lacks sufficient reason to do, a possibility that is anyway familiar from cases of akrasia. The problem arises in this case only because one cannot expect one's intention to be stable over time when one foresees that one will believe that one has compelling reasons to reconsider and that one will be sufficiently rational to respond accordingly. This problem would not arise if the proposal were to intend to drink the toxin not tomorrow, but instead right now. Most people could, I imagine, form this intention. Nor would it arise if one foresaw that one would stick to the intention with irrational resoluteness, either by refusing to reconsider one's intention, or by refusing akratically to respond to reconsideration.

are evidential considerations, and that subjects in practice cannot but treat them as identical. I offer an argument for this, albeit for a different purpose, in section 3. This means that my account of the relationship between B+ and B−, on the one hand, and C+ and C−, on the other, must take a different form from my account of the relationship between I+ and I− and C+ and C−. Giving an account of the relationship between B+ and B− and C+ and C− is part of the unfinished business that I catalogue, but don't undertake, at the end of this paper.

Let me finish this discussion by addressing some final, stray worries about certain of these requirements. First, one might object, against B+, that there cannot be a rational requirement to believe that *p* if one believes that there is conclusive evidence that *p*, because this requirement cannot be violated.¹⁹ To believe that there is conclusive evidence that *p*, one might claim, just is to believe that *p*. But this is not obviously so. While it is difficult to spell them out precisely, there are other marks of believing that *p*, besides believing that there is evidence that *p*. For example, there is the mark of being prepared to rely on its being the case that *p* in settling theoretical or practical questions that one believes depend on whether it is the case that *p*, at least in so far as one seeks to settle those questions. One might lack enough of these other marks to be correctly described as not believing that *p*, even though one believes that there is conclusive evidence that *p*.

Second, one might have doubts about I+ in light of the argument of Arpaly (2000) that it is sometimes 'more rational' to act against one's judgement about what one ought to do than to act with it. Arpaly's argument, however, has little bearing on I+, understood as a *local* requirement of *subjective* rationality. To begin with, Arpaly's claim is that in certain cases, acting against one's judgement is more rational *overall*, or *on balance*, than acting with it, given other features of one's psychology and situation. Whatever this means, it is not obviously incompatible with the claim that, in so doing, one violates a local requirement of rationality. It might mean, for example, that one would have violated even more local requirements otherwise. As Smith (2004, p. 190) observes, Arpaly's agents appear to purchase their 'global rationality at the cost of local irrationality.' More importantly, Arpaly appears to be discussing *objective*, rather than *subjective*, rationality, although she herself does not draw this distinction. The sense in which Arpaly's agents are 'more rational' in acting against their judgement is that in so doing they make the choice that the reasons actually support, or the

¹⁹ Moreover, there are questions about the premiss on which this objection implicitly relies: that it must be possible to violate normative requirements. See Lavin (2004).

choice that the evidence actually indicates that the reasons support. I discuss related issues at greater length in subsection 1.8.

Finally, one might worry that I⁻ is too austere, given that there are cases in which it is not irrational to do X, even though one believes that one has no reason to do it.²⁰ But I⁻ is compatible with such cases. In such cases, one also believes that one does not need any reason to X. Thus, even if one believes that one does not have a reason to X, one still believes that, trivially, one has *sufficient* reason to X. (If I write a cheque for zero dollars and zero cents, then even if I do not have funds in my account, I still have sufficient funds to cover the cheque.) If one believed that one lacked sufficient reason to X—if one believed that one did need reason to X in the circumstances, but had no such reason—then it would be irrational of one to intend to X.

1.4 First argument that wide-scope requirements fail the Reasoning Test: There must be a content to reason from

Now back to the main thread of argument. In this subsection and the next, I argue that wide-scope versions of B⁺, B⁻, I⁺, and I⁻ fail the Reasoning Test. It follows that these rational requirements have narrow scope. If any rational requirements have narrow scope, then it cannot be the case that we have conclusive reasons to comply with rational requirements in general. For if we had conclusive reasons to comply with these narrow-scope requirements, then there would be bootstrapping.

Recall I⁺. In explicitly wide-scope form, it would be:

I+WS: Rationality requires one (either not to believe that one has conclusive reason to X, or to intend to X).

I+WS says that, when one is conflicted in this way—that is, when one believes that one has conclusive reason to X, but does not intend to X—there are two ways to resolve this conflict rationally. Either one can form the intention on the basis of the content of one's belief about one's reasons, or one can revise one's assessment of one's reasons on the basis of the content of one's not intending to X. In terms of the Reasoning Test, I+WS is correct only if (i) one can reason from the content of one's belief that one has conclusive reason to X to intending to X and (ii) one can reason from the content of one's not intending to X to a revision of one's belief that one has conclusive reason to X.

One certainly can reason from the content of one's belief that one has conclusive reason to X to an intention to X. So one can rationally

²⁰ A point pressed separately by Broome, David Kaplan, and John Searle.

resolve the conflict in that way. But one cannot reason from the content of one's not intending to *X* to a revision of one's belief that one has conclusive reason to *X*. One cannot reason from the content of one's not intending to *X* to anything at all, *because there is no such content*. Not intending to *X* is simply lacking an attitude. The *lack* of an attitude has no content. (Intending not to *X* does have content, namely, that one will not *X*. But intending *not* to *X* is not the same as *not* intending to *X*.) Of course, starting from such a conflict-state, one might simply lose one's belief that one has conclusive reason to *X*. That can happen, after all. And if it happens, then, in some sense, the conflict is resolved; one is no longer in the conflict-state. But losing one's belief on no basis at all would not be *rationally resolving* the conflict. So I+ has narrow scope:

I+NS: If one believes that one has conclusive reason to *X*, then rationality requires one to intend to *X*.

This reflects the fact that there is only one way for one to resolve this conflict rationally. There is only one direction for one's reasoning to take. A similar argument shows that B+ has narrow scope.

1.5 Second argument that wide-scope requirements fail the Reasoning Test: Reasoning is 'downstream'

The same argument cannot be used to show that B− and I− have narrow scope. For they govern not conflicts between an attitude and the lack of an attitude, but instead conflicts between two attitudes. A different argument is needed.

Consider I−. In explicitly wide-scope form, it would be:

I−WS: Rationality requires one (either not to believe that one lacks sufficient reason to *X*, or not to intend to *X*).

I−WS says that, when one is conflicted in this way—that is, when one believes that one lacks sufficient reason to *X*, but still intends to *X*—there are two ways to resolve this conflict rationally. Either one can drop the intention on the basis of the content of one's belief about one's reasons, or one can revise one's assessment of one's reasons on the basis of the content of one's intention. In terms of the Reasoning Test, I−WS is correct only if (i) one can reason from the content of one's belief that one lacks sufficient reason to *X* to dropping one's intention to *X* and (ii) one can reason from the content of one's intending to *X* to revising one's belief that one lacks sufficient reason to *X*.

One can reason from the content of one's belief that one lacks sufficient reason to *X* to dropping one's intention to *X*. One can rationally

resolve the conflict in this way. But one cannot reason from the content of one's intention to *X* to revising one's belief that one lacks sufficient reason to *X*. It is not reasoning to cling to what one judges to be an unfounded intention and to support it by revising one's belief about one's reasons. It is a kind of wishful thinking or self-deception. Consider, to a first approximation, how one would express this transition to oneself: 'The facts of my situation do not give me sufficient reason to *X*. I hereby commit to doing *X*. Thus, all along, the facts of my situation gave me sufficient reason to *X*.'²¹ I say 'to a first approximation,' because the 'thus' seems out of place. What relation could one be taking it to express?

The rational requirement governing such conflicts thus has narrow scope:

I–NS: If one believes that one lacks sufficient reason to *X*, then rationality requires one not to intend to *X*.²²

This reflects the fact that there is only one way for one to resolve this conflict rationally. There is only one direction for one's reasoning to take. This is not because, as was the problem with I+, one has no other content to reason from; the intention to *X* has a content. It is instead because to reason is to be guided by one's assessment of one's reasons. One can't reason 'upstream'—from one's attitudes to a reassessment of one's reasons for them—only 'downstream'—from one's assessment of one's reasons for one's attitudes to the formation, retention, or revision of those attitudes. Rationality requires one, in forming, retaining, or revising one's attitudes, to follow the downstream current. The point is

²¹ Some might believe that intending to *X* gives one reason to *X*; willing is the source of reasons. But I take it that nonreductionists, to whom the present argument is addressed, do not believe this. What provides one with reasons, according to nonreductionists, are the facts of one's situation (which only in special cases will include the fact that one intends something). The phrases 'all along' and 'facts of my situation' are meant to bring this out.

²² Glenn Ross raises the following objection against the narrow-scope readings of I+ and I–. If someone, at one and the same time, believes that she has conclusive reason to *X* and believes that she lacks sufficient reason to *X*, then I+ and I– will require incompatible things of her. She will be required, at one and the same time, to intend to *X* and not to intend to *X*. Whatever she does next, she will be in some way irrational. This seems to me, however, the right thing to say about such a case. She has backed herself into a corner. So it is unclear, on reflection, what the objection is supposed to be. Why can't it be the case that rationality requires one to intend to *X* and not to intend to *X*? Because 'ought' of reason implies 'can'? But it would beg the question, at this point, to assume that the 'ought' of rationality is the 'ought' of reason. If the 'ought' of rationality is understood along the lines of the Transparency Account, then it is entirely explicable how one could be under conflicting rational requirements. If someone believes that she has conclusive reason for an attitude, then, as it seems to her, she ought to have it. And if she also believes that she lacks sufficient reason for that attitude, then, as it seems to her, she also ought not to have it. She feels two conflicting normative 'pressures,' and whatever she does next, she will be resisting one of them.

not that one shouldn't reason upstream, that progressing upstream is poor reasoning, but that one simply cannot reason upstream, that progressing upstream is not recognizable as reasoning at all. It is some other process, such as self-deception or wishful thinking. A parallel argument shows that B- has narrow scope.

1.6 First objection, to both arguments: The examples are underdescribed

One's initial response to the question, 'How may one rationally respond to the conflict of, say, believing that one lacks sufficient reason to *X*, but intending to *X*?' might be to suspend judgement. 'It all depends,' one might think, 'on the facts of the case. If *X*-ing is, say, watching mindless television late into the night before an important, early-morning job interview, then, yes, the rational thing to do is to drop one's intention. But if *X*-ing is watching television to learn the latest status of a threatening tornado, then maybe dropping one's intention is not the rational thing. Perhaps the rational thing is to revise one's belief that one lacks sufficient reason to *X*.'

Although this request to hear the facts of the case is perfectly judicious in other contexts, it is misplaced here. Our topic is *subjective* rationality, which is a matter of the relations among one's attitudes *in abstraction from the reasons for them*: that is, in abstraction from the facts of one's circumstances that might actually favour a given attitude. We should not need any information about the facts of one's circumstances in order to judge whether one is responding to the conflict in a *subjectively* rational way. If our judgements vary with such information, then we must be focusing on the wrong thing; we must be making judgements about something other than subjective rationality. There is, as I have granted, a notion of *objective* rationality, according to which the 'rational thing to believe' is whatever the evidence *in fact* supports and the 'rational thing to do' is whatever the features of one's situation *actually* weigh in favour of doing. Certainly, our judgements about the 'rational thing to do,' understood in this objective sense, *will* vary with the facts of the case. But this is not the kind of rationality with which we are concerned.

It might be worthwhile, at this juncture, to pause to get our methodological bearings. Our topic, as I have said, is subjective rationality, which is a matter of the relations among one's attitudes, viewed in abstraction from the reasons for them. Subjective rationality is, roughly speaking, a matter of maintaining consistency among one's attitudes. It is what is violated by such conflicts as willing the end, but not the apparent means; believing *p* and *p* entails not-*q*, while also believing *q*;

and believing that one has conclusive reason to *X*, but not intending to *X*. The question we are asking is: In what sense 'ought' we to be subjectively rational, to be consistent? In particular: Do we have *reasons* to be subjectively rational, as Broome suggests?

Now, one might worry that my definition of subjective rationality prejudices the answer. 'If subjective rationality is defined as a matter of the relations among one's attitudes, viewed in abstraction from the reasons for them, then doesn't it follow, by definition, that there are no reasons to be subjectively rational?' When I say that subjective rationality is a matter of the relations among one's attitudes, viewed in abstraction from the reasons for them, I mean that what subjective rationality *consists in*, what the *content* of its requirements are, is independent of any reasons that one in fact has for particular attitudes. We can settle which processes of attitude formation, retention, and revision are subjectively rational without knowing whether there are in fact good reasons for any of the attitudes involved. I take it that we have an implicit sense of what counts as subjective rationality, or consistency, which we can elicit by considering particular cases. Having identified what subjective rationality is, what the content of its requirements are, we can then ask whether there are *reasons*, in general, to be subjectively rational, whether we *ought* to comply with those requirements. The answer to this question is not prejudiced—although, of course, this paper tries to argue that the answer is 'no.' In any event, this methodology—first determine the content of the requirements of subjective rationality without invoking any substantive claims about reasons, then explain in what sense those requirements are normative—is not controversial. It is accepted by all of the participants to the debate whom I have mentioned.

Once one has clearly distinguished subjective rationality from objective rationality, however, one might begin to have doubts about the importance of the former. 'What is so rational about subjective rationality?' one might wonder; 'What is so rational about consistency?' What I take these questions to ask is: 'What is so objectively rational about subjective rationality? Why think that we have reason, in general, to be consistent?' My reply is: 'Good question, and precisely the question we are asking.' My overarching argument, indeed, is that it is *not* objectively rational, in general, to be subjectively rational. We do not have *reason*, in general, to be consistent. The sense in which we 'ought' to be consistent—the sense in which subjective rationality is 'normative'—must be understood in a different way. The Transparency Account attempts to offer such an alternative understanding.

1.7 Second objection, to both arguments: It is possible to rationally exit conflicts by revising the upstream attitude

Having made these methodological remarks, let me turn to a second objection to my two arguments that B+, I+, B−, and I− have narrow scope. I have been arguing that when one faces one of the conflicts that B+, I+, B−, and I− govern, there is only one rational way for one to resolve that conflict. For example, when one believes that one has conclusive reason to X but does not intend to X, the only way for one to resolve the conflict rationally is for one to form the intention to X. So, I have claimed, I+ has narrow scope. ‘But what if one has a belief, further upstream, that there is *insufficient evidence* that one has conclusive reason to X?’ one might ask. ‘Then surely it would be rational of one to revise, on that basis, one’s downstream belief that one has conclusive reason to X. In doing so, one would *leave* the conflict state of believing that one has conclusive reason to X and not intending to X. And one would leave that conflict-state *rationally*. Therefore, one *can* rationally resolve the conflict by revising one’s belief that one has conclusive reason to X. So I+ has wide scope after all.’

It is true that in revising one’s belief that one has conclusive reason to X, one would be *rationally responding to some conflict*. And it is true that in revising one’s belief that one has conclusive reason to X, one would *leave the conflict that I+ governs*: the conflict of believing that one has conclusive reason to X but not intending to X. However, in revising one’s belief that one has conclusive reason to X, one would not be *rationally responding to the conflict that I+ governs*. Instead, one would be rationally responding to the conflict, further upstream, of believing that there is not sufficient evidence that one has conclusive reason to X, but believing that one has conclusive reason to X. One would not be reasoning, *per impossible*, from the content of one’s *not intending* to X to a revision of one’s belief that one has conclusive reason to X. One would be reasoning, instead, from the content of one’s *upstream belief* that there is not sufficient evidence that one has conclusive reason to X to a revision of one’s belief that one has conclusive reason to X. One would be complying not with the requirement that governs the conflict of believing that one has conclusive reason to X, but not intending to X, but instead with the requirement that governs the conflict of believing that there is not sufficient evidence that *p*, but believing that *p*. In other words, one would be complying not with I+, but instead with:

B−NS: If one believes that there is insufficient evidence that *p*, then one is rationally required not to believe that *p*.

In sum, the fact that one may rationally revise one's downstream belief that one has conclusive reason to X , on the basis of the content of an upstream belief that there is not sufficient evidence that one has conclusive reason to X , is entirely in keeping with the point that reasoning follows the downstream current. It does nothing to undermine the claim that $I+$ has narrow scope: that $I+NS$ correctly indicates how one may rationally respond to the conflict of believing that one has conclusive reason to X , but not intending to X . It only reflects the fact that there is *another* narrow-scope requirement, $B-NS$, which indicates how one may rationally respond to the conflict of believing that there is insufficient evidence that p , but believing that p .

One might worry, however, about what this analysis implies in a case in which I believe that there is insufficient evidence that I have conclusive reason to X , believe that I have conclusive reason to X , and do not intend to X . $B-NS$ implies that I am rationally required to revise my belief that I have conclusive reason to X , and $I+NS$ implies that I am rationally required to intend to X . 'This seems odd,' one might say. 'In this case, I am rationally required only to revise my belief about my reasons, not also to have the intention that the belief to be revised recommends. Somehow these requirements—or, rather, *prima facie* requirements—interact, with one cancelling the other out. So $I+NS$ is not valid in all instances. And this shows that the bootstrapping argument, which relies on $I+NS$, cannot be made.'

It is not clear why this claim about the interaction of rational requirements, if it were correct, should threaten the bootstrapping argument. All that the argument needs is one case in which $I+NS$ is valid. And the objection grants that $I+NS$ would be valid in a case in which one believes that there *is* sufficient evidence that one has conclusive reason to X . But, in any event, I do not think that the claim is correct. Rational requirements do not interact in this way. They are 'local,' in the sense in which I stressed earlier. One can be rationally required, by a given principle, to revise a belief, while at the same time being rationally required, by another principle, to have an intention.²³

²³The Transparency Account makes sense of this. In so far as one believes that there is not sufficient evidence that one has conclusive reason to X , one will feel normative 'pressure' to revise one's belief that one has conclusive reason to X . And in so far as one believes that one has conclusive reason to X , one will feel normative 'pressure' to intend to X .

1.8 *Third objection, to the second argument: Apparent examples of upstream reasoning*

Finally, we come to a particularly formidable objection, which targets my second argument. Although the objection does not bear on my case that I+ and B+ have narrow scope, it does bear on my case that I− and B− have narrow scope, for it claims that there are cases of upstream reasoning. ‘Suppose that I believe that I lack sufficient reason to X, but still intend to X. Suppose, further, that I believe, from experience, that whenever I have a recalcitrant, akratic intention of this kind, which defies my assessment of my reasons, the intention turns out to be right. That is, it turns out that my belief was false and that I did have sufficient reason to X. Surely, in this case, I can revise my belief that I lack sufficient reason to X in light of my intention to X. Isn’t this a case of upstream reasoning? Notice that there is nothing bizarre or contrived about this case. Often our funny feelings, gut instincts, and such like rebel against our “better” judgement. When this happens, there is nothing necessarily irrational about trusting these “downstream” attitudes and revising our “upstream” attitudes in light of them.’

In order to answer this objection, we need to be precise about the case that we are considering. Let me clear the way by considering two simple cases. Later, I will address a more complicated case, which involves a kind of divided psychology, in which one’s beliefs are not accessible to conscious reflection.

Suppose—to consider the first simple case—that one is reasoning from

(A1) the content of one’s belief that one intends to X

and

(A2) the content of one’s belief that if one intends to X, then that is evidence that one has sufficient reason to X,

to

(A3) a revision of one’s belief that one lacks sufficient reason to X.

This is not upstream reasoning. It is downstream reasoning, from the content of the *upstream* belief that there is evidence that one has sufficient reason to X to a revision of one’s *downstream* belief that one lacks sufficient reason to X. As we observed in reply to the previous objection, I−NS does not rule out revising one’s belief that one lacks sufficient reason to X on the basis of the content of one’s upstream

assessment of the evidence for that belief. (As we saw, I–NS has nothing to say about such revisions. They are governed instead by B–NS.)

What I–NS rules out, and what would be upstream reasoning, is when, without any beliefs about the evidence, conscious or otherwise, one transitions directly from

(B1) the content of one's intention to *X*

to

(A3) a revision of one's belief that one lacks sufficient reason to *X*.

This is the second simple case. One might express this transition to oneself by: 'The facts of my situation do not give me reason to *X*. I hereby commit to doing *X*. All along, the facts of my situation gave me reason to *X*.' This transition is simply not recognizable as reasoning. Such a transition might occur, of course; one's intention might directly cause a revision of one's belief. But this would be some process other than reasoning, such as wishful thinking or self-deception.

The transition from (A1) and (A2) to (A3) can easily be confused with the transition from (B1) to (A3), because much the same words can be used to describe them. In the first case, one reasons from the content of one's *beliefs that* one intends to *X* and that that intention indicates that one has sufficient reason to *X*. This might be described as 'reasoning from one's intention to *X*.' So described, it may sound like the same thing as reasoning from the content of one's *intention* to *X*, by itself, which is what one would be reasoning from in the second case, were it in fact a case of reasoning. But the two cases are quite different.

To distinguish the two cases, it may help to observe that, in the first case, one's intention functions simply as a fact, or apparent fact, about one's circumstances: a fact, or apparent fact, that one takes to indicate that one has sufficient reason to *X*. Thus, one's reasoning from (A1) and (A2) to (A3) has the same form as reasoning from:

(C1) the content of one's belief that the town sage is telling one to *X*

and

(C2) the content of one's belief that if the town sage is telling one to *X*, then that is evidence that one has sufficient reason to *X*,

to

(A3) a revision of one's belief that one lacks sufficient reason to *X*.

This is clearly downstream reasoning.

Put another way, in reasoning from (A₁) and (A₂) to (A₃), one is not, as it were, reasoning ‘from inside’ one’s intention, as one would be in reasoning from (B₁) to (A₃), were it a case of reasoning (and, arguably, as one reasons, in instrumental cases, from intending the end to intending the apparent means). Instead, one is reasoning ‘from outside’ one’s intention. One views one’s intending to *X* as a fact that one is confronted with, and one asks what that fact indicates about one’s situation. Notice that unless one views one’s intention as something that one is confronted with—as something that eludes one’s control, either by arising unbidden, or by resisting alteration—one cannot see it as evidence of what one ought to intend. If, by contrast, one took intending to *X* to be entirely within one’s control—if one inhabited one’s intention in such a way that one might express it as, ‘I hereby commit to doing *X*’—then it is unclear how one could view it as providing evidence that one ought to intend to *X*.

Still, there may be some temptation to think that it is rational for one to transition from (B₁) to (A₃), as one does in the second case, a temptation that results from a confusion of objective and subjective rationality. Suppose that one intends to *X*, but one hasn’t reflected on this fact. And suppose that one’s intending to *X* is in fact evidence that one has sufficient reason to *X*, even though one does not know that it is. In this case, one cannot reason from (A₁) and (A₂) to (A₃), because one has not reflected on (A₁) and does not believe (A₂). Nevertheless, suppose that one transitions from (B₁) to (A₃): that is, let us suppose, one’s intention to *X* directly causes one’s belief that one lacks sufficient reason to *X* to change, without the presence of any thoughts about the evidence for or against that belief. ‘It is rational to transition in this way, even if one does not realize that it is so,’ one might argue. ‘For it results in one’s having the belief that the evidence in fact supports. Moreover, a disposition to transition in this way would reliably lead one to have beliefs that the evidence supported. Whether or not one realized it, such a disposition would be something to welcome.’

By ‘rational,’ it seems, the objector means ‘objectively rational.’ The evidence, although one is not aware of it, in fact argues in favour of revising one’s belief. So the *objectively* rational thing to do is to revise it. Likewise, a disposition to transition from one’s recalcitrant, akratic intentions to *X* to revisions of one’s beliefs that one lacks sufficient reason to *X* would be a disposition that reliably led one to hold beliefs that the evidence in fact supported. So it would be disposition that led one to hold beliefs that it was *objectively* rational to hold. A clear indication

that the judgements here are judgements of objective rationality is the fact that they are sensitive to the reasons—in this instance, to the evidentiary reasons—that actually obtain in the case. If we supposed that one's recalcitrant, akratic intentions to *X* were not evidence—did not indicate—that one had sufficient reason to *X*, then we would not judge that one's transitioning from (B₁) to (A₃) was rational in this sense (or, indeed, in any other sense).

The point to which this objection appeals is deep, important, and (outside of the philosophy of the emotions) largely overlooked. Being led by our gut instincts may be a more reliable path to objectively rational belief and action than our consciously evaluating reasons. That is, being led by our instincts may more reliably result in our having the beliefs that the evidence in fact supports and in our doing what we in fact have reason to do. But, to repeat a point made earlier, the present question is whether the transition from (B₁) to (A₃) is *subjectively* rational, which is a matter of the relations among one's attitudes, *in abstraction from the reasons for them*. In fact, provided that the distinction between objective and subjective rationality is clearly drawn, this point, that we would often do better to follow our instincts, is entirely friendly. What it suggests is that we do not have reason to be subjectively rational, since being subjectively rational may not lead us to have the attitudes that we have reason to have. In the end, this is precisely what I am trying to argue.

Having addressed the two simple cases, let us consider a more complicated one. Suppose that one's psychology is divided. At the level of conscious reflection, one believes that one lacks sufficient reason to *X*, and one sees no evidence to the contrary. But at a different level, not transparent to conscious reflection, one believes that there is (rather decisive) evidence to the contrary. Although this belief does not percolate into consciousness, it exercises its influence by reinforcing one's akratic intention. Suppose that this intention, backed by this hidden belief about the evidence, causes a revision of one's belief that one lacks sufficient reason to *X*. In this way, one transitions from (B₁) to (A₃). To make sure that our judgement about this case is not influenced by considerations of objective rationality, let us abstract away from the facts of the case; one's hidden belief may be on the right track, or it may not. Even so, this transition from (B₁) to (A₃) may strike us as at least a *better* candidate for subjective rationality than the same transition in the second simple case, in which one had *no* belief, at *any* level, that there was evidence against one's belief that one lacked sufficient reason to *X*. More precisely, I suspect that we will be ambiva-

lent about whether one is being subjectively rational in this case. On the one hand, we are more inclined to count it as subjectively rational than the transition from (B₁) to (A₃) in the second simple case, in which one had no beliefs, at any level, about the evidence against the belief that one lacked sufficient reason to *X*. On the other hand, we are less inclined to count it as rational than the transition from (A₁) and (A₂) to (A₃) in the first simple case, in which one's beliefs about the evidence were transparently available to one. The present case is not quite as clear cut as that.

The question for my account is whether it explains this ambivalence. I suggest that it does. The ambivalence stems, I think, from our deeper ambivalence about how, in cases in which one's psychology is divided, the boundaries of 'one's psychology' are to be drawn. On the one hand, it seems distorting to identify 'one's psychology' with what is available to conscious reflection, since so much is going on beneath the surface. On the other hand, it seems distorting to identify 'one's psychology' with the whole, since so much is hidden from one. To the extent that we identify 'one's psychology' strictly with *what is present to reflective consciousness*, then all that occurs within 'one's psychology' is the bare transition from (B₁) to (A₃). Thus, as my account would predict, we are less inclined to count it as subjectively rational. To the extent that we identify 'one's psychology' with *all of one's attitudes, present at any level*, then what is occurring within 'one's psychology' is closer to one's revising one's belief that one lacks sufficient reason to *X* on the basis of the content of one's belief that there is evidence that one has sufficient reason to *X*. (After all, one believes that there is evidence that one has sufficient reason to *X*, and it is only because of this belief that one revises one's belief that one lacks sufficient reason to *X*.) Revising one's belief in this way is downstream reasoning, and thus, as my account would predict, we are more inclined to count it as subjectively rational. In sum, since we are ambivalent about which to identify 'one's psychology' with, we are, as my account would predict, ambivalent about whether this transition is subjectively rational.

Put another way, my account of the content of rational requirements implies that in cases of divided psychology such as this, there will be no clear answer what (subjective) rationality requires. To the extent that we identify one's psychology with what is present to reflective consciousness, then I–NS applies, and rationality requires one to revise one's intention to *X*. To the extent that we identify one's psychology with what is present at any level, then B–NS applies, and rationality requires one to revise one's belief that one lacks sufficient reason to *X*.

As I have said, this seems to me the right result. We are, in fact, ambivalent about such cases.²⁴

Of course, once we introduce the facts of the case, it may seem obvious to us what ‘rationality’ requires. If we know that one’s recalcitrant intentions reliably indicate what one ought to do, then we will think it clear that ‘rationality’ requires one to revise one’s belief that one lacks sufficient reason to *X*. No ambivalence there! But, once again, we are understanding ‘rationality’ here as ‘objective rationality.’ Objective rationality requires one to believe what the evidence in fact supports. Since we know that the evidence in fact supports revising one’s belief that one lacks sufficient reason to *X*, we judge that objective rationality requires one to revise one’s belief. However, to repeat a point made several times now, our concern is with subjective rationality.

1.9 Conclusion: We cannot have conclusive reason to comply with rational requirements generally

If this reply is successful, then the argument that B– and I– have narrow scope stands. Even if it is not successful, the argument that B+ and I+ have narrow scope is unaffected. In either case, it follows that we do not have reasons to comply with rational requirements in general.

Suppose, on the contrary, that we had, as Broome once thought, conclusive reason to comply with rational requirements in general. That is, suppose we accept the:

Reasons Claim: If one is rationally required to *Z*, then one has conclusive²⁵ reason to *Z*,

²⁴ The Transparency Account coheres with the view that it is indeterminate what (subjective) rationality requires of one in this case. When one’s psychology is divided, ‘one’ will feel conflicting normative pressures. To the extent that ‘one’—one’s conscious, reflective self—believes that one lacks sufficient reason to *X*, ‘one’ will feel normative pressure to revise one’s intention to *X*. But to the extent that ‘one’—one’s unconscious, intuitively perceptive self—believes that there is insufficient evidence that one lacks sufficient reason to *X*, ‘one’ will feel normative pressure to revise one’s belief that one lacks sufficient reason to *X*. Likewise, to the extent that we take ourselves to be addressing one’s conscious, reflective self—as we would be if we said something out loud to one here and now—then our advice (or quasi-advice) will be that one ‘ought rationally’ to revise one’s intention to *X*. To the extent that we take ourselves to be addressing one’s whole self, including one’s unconscious attitudes—which might be analogous to the third-person cases that I discuss in section 5—then our advice (or quasi-advice) will be that one ‘ought rationally’ to revise one’s belief that one lacks sufficient reason to *X*.

²⁵ Why conclusive reason, instead of simply *pro tanto* reason? Because it seems necessary to capture the way in which rational requirements correspond to ‘oughts’: what Broome (1999) calls their ‘strictness.’ In spite of this, one might be attracted to ‘slackening’ the Reasons Claim to a claim about *pro tanto* reasons, as a way of avoiding the bootstrapping result that follows: namely, that if one believes that one lacks sufficient reason to *X*, then one has conclusive reason not to in-

where ‘to *Z*’ might be filled in by ‘to have some attitude, *A*,’ ‘not to have some attitude, *A*,’ or ‘to ensure that whatever attitudes one has stand in some relation, *R*.’ The Reasons Claim is fine when applied to I–WS, since it entails only:

One has conclusive reason (either not to believe that one lacks sufficient reason to *X*, or not to intend to *X*).

But, as I have been arguing, I– has narrow scope; I–NS is its proper formulation. When applied to I–NS, the Reasons Claim implies:

If one believes that one lacks sufficient reason to *X*, then one has conclusive reason not to intend to *X*.

Not all beliefs would be self-justifying, but beliefs about lacking sufficient reasons for action would be. This result is no more acceptable. We must conclude, therefore, that I–NS represents at least one rational requirement with which we do not have conclusive reason to comply. This is enough to defeat the Reasons Claim.

It seems that a theory of error ought to be forthcoming. For a long time after reading Broome, I was convinced by his arguments that rational requirements have wide scope. I now believe that some rational requirements have narrow scope. Why did I overlook these? Part of the reason is that Broome takes rational requirements to be state-requirements, and state-requirements have wide scope. For example, the state of not believing that one lacks sufficient reason to *X* and intending to *X* satisfies, just as well as the state of believing that one lacks sufficient reason to *X* and not intending to *X*, the ban on the conflict-state of believing that one lacks sufficient reason to *X* and intending to *X*. In other words, if you come on the scene, knowing only that, as things now stand, I am in the state of not believing that I lack sufficient reason to *X* and intending to *X*, then you have no grounds to accuse me of irrationality. It is only when you learn that I got there by revising my assessment of my reasons to accord with my intention that you come to have such grounds. What you learn, in that case, is that I violated the *process*-requirement governing how to *resolve* the conflict-state of believing that I lack sufficient reason to *X* and intending to *X*. It is only once we turn from state-requirements to process-requirements—only once we

tend to *X*. But the bootstrapping result that one even has *pro tanto* reason not to intend to *X* seems unacceptable on much the same grounds. Imagine some *X* that one has no reason at all not to intend to do. Do we now want to say that one’s believing that one lacks sufficient reason to *X* makes it the case that one does have such a reason, which carries some weight, although perhaps not decisive, against all of the others?

shift our focus from the rationality of synchronic states to the rationality of diachronic transitions among them—that we come to see that some rational requirements have narrow scope.

The other part of the reason, I think, is that Broome focuses largely on rational requirements of logical consistency and means-end coherence. These requirements, even when construed as process requirements, are not simple narrow-scope requirements. For example, suppose that one is in the conflict-state of believing that p and believing that q , against (to simplify things) the background of a fixed belief that p entails not- q . Is it plausible to claim that there is only one way for one to resolve this conflict rationally? Surely not. One might resolve it by reasoning in accordance with *modus ponens*, thus revising one's belief that q . And one might resolve it by reasoning in accordance with *modus tollens*, thus revising one's belief that p . So the relevant rational requirement is clearly not a simple narrow-scope requirement. After all, no unique response is required; one might revise one belief, or the other, or both.

While this conflict is not governed by a simple narrow-scope requirement, however, it is not governed by a wide-scope requirement either. It is governed instead by a *disjunction* of narrow-scope requirements. Notice that as soon as we bring into view one's upstream beliefs about the evidence that p and that q , it is no longer clear that one has more than one way of rationally resolving the conflict. Suppose that one judges that there is stronger evidence that p than there is that q . Then it seems that one is rationally required to revise one's belief that q . Suppose that one judges that there is weaker evidence that p than there is that q . Then one is rationally required to revise one's belief that p . Finally, suppose that one judges that there is neither stronger, nor weaker evidence that p than there is that q . Then one is rationally required to suspend judgement: that is, to revise both beliefs. In sum, once one's upstream beliefs about the evidence are taken into account, those beliefs determine a uniquely rational way to resolve the conflict.²⁶ There are, in effect, three *narrow-scope* requirements at work here: (1) if one believes that the evidence that p is stronger than the evidence that q , then one is rationally required not to believe that q ; (2) if one believes that the evidence that p is weaker than the evidence that q , then one is rationally required not to believe that p ; and (3) if one believes that the

²⁶ Broome (2003b, lec. 2 p. 10), says something similar: 'I expect your higher-order normative beliefs—about what you ought to believe or about what rationality requires you to believe—will influence the direction of your reasoning'. However, Broome does not say that rationality requires that one reason in a particular direction. Note also that his understanding of 'reasoning' is more restricted than mine.

evidence that p is neither stronger, nor weaker than the evidence that q , then one is rationally required not to believe that p and not to believe that q . If we don't know what one thinks about the evidence, then we do not know which narrow-scope requirement one is under. All we can say is that (1) either one is rationally required to revise one's belief that q , (2) or one is rationally required to revise one's belief that p , or (3) or one is rationally required to revise both beliefs. *For all we know, modus ponens* might be the rational response, or *modus tollens*, or suspension of belief. But this does not mean that all of these are *in fact* rational responses, as a wide-scope requirement would suggest.²⁷

This helps to confirm a conjecture that I make in conjunction with the Transparency Account. In section 5, I suggest that all rational requirements can be derived from two 'core' narrow-scope requirements. One might find it puzzling how a requirement like the one just discussed, which seems to permit more than one kind of resolution, could be derived from narrow-scope requirements. The foregoing discussion illustrates one such derivation.

2. Second argument that there are no reasons to be rational: what would the reasons be?

The bootstrapping problem is a reason to doubt that there are reasons to comply with certain rational requirements, namely B+, I+, B−, and I−. But there is a ground for doubting that there are reasons to comply with *any* rational requirements. If there were reasons to comply with rational requirements, what would they be? We can contrive situations, of course, in which people have instrumental reasons for avoiding irrationality. Caligula might threaten you with a life of torment unless you conform to I+. (This would give literal meaning to the phrase, 'on pain

²⁷ It is easy to confuse a disjunction of requirements of attitudes—that is, this disjunction of narrow-scope requirements—with a requirement of a disjunction of attitudes—that is, a wide-scope requirement. But the two are different, as an example may help to illustrate. Suppose that I moved from California to Massachusetts on January 2, 2004. I know that I need to file one state tax return as a resident and one as a nonresident, but I don't know which. So I call the help desk to find out whether for tax purposes residency is a matter of where one started, or where one spent more time. What I know, as I wait on hold, is that either I am legally required to file as a resident of California and not as a resident of Massachusetts, or that I am legally required to file as a resident of Massachusetts and not as a resident of California. I do not know, because it isn't true, that I am legally required (either to file as a resident of California and not as a resident of Massachusetts, or to file as a resident of Massachusetts and not as a resident of California). This requirement would be satisfied by my filing as a resident of California and not a resident of Massachusetts. But there is no legal requirement that is satisfied by my filing as a resident of California and not as a resident of Massachusetts. As the tax code works, what matters is where one spent more time. I am legally required to file as a resident of Massachusetts and not as a resident of California, and doing anything else would be against the law.

of irrationality.) In this case, you would have a conclusive reason, of a familiar kind, to comply. The normativity of rationality, however, cannot be explained by reasons of this kind. Rationality is normative for any believer or intender, no matter what his circumstances. If its normativity consists in reason to conform to rational requirements, then this reason must be more general.

One might suggest the following, more general reason: that by conforming to rational requirements, one is more likely to believe and do what one ought. Like the reason that Caligula gives you, this reason would be instrumental. One would not comply with rational requirements for their own sake, but because doing so brought about something for which one had independent reasons. Nevertheless, it might be argued, one *always* has these reasons, so the justification is general. It does not depend on contrived circumstances.

The problem is that it is not true, in any given case, that complying with rational requirements leads one to believe and do what one ought. Suppose I believe that I lack sufficient reason to believe that p , but believe that p . As things stand, I violate B-. Then I comply by not believing that p . It might be the case, however, that I have reason to believe that p . So being rational leads me to lose a belief that I have reason to have. To this, the answer is bound to come that while rationality may lead one astray in any given case, it is the best policy in the long run. If one complies with rational requirements as a rule, then over the long run one is more likely to believe and do what one has reason to.

It is not clear why this should be so. The net result of revising my attitudes in accordance with rational requirements might be to adopt many attitudes for which I have no reason, and to abandon many attitudes for which I have.²⁸ In any event, even if it is true that we have this reason to comply with rational requirements as a rule, it does not follow that we have this reason to comply with them in any particular case. Yet the 'ought' of rationality applies in each particular case. When we say that someone 'ought rationally' to have an attitude, we are saying something about what 'ought' to happen here and now. This problem is parallel to the one that undoes traditional forms of rule utilitarianism. Why is there reason to follow, in this particular case, the rule that promises utility over the long run, if violating it, in this particular case, promises even more utility? Likewise, why is there reason to

²⁸ Contrast conforming to norms of *objective* epistemic rationality, such as, perhaps: If there is reason to believe that p , and there is reason to believe that p entails q , then there is reason to believe q . Conforming to such norms will lead one to have beliefs that one has reason to have. The subjective rational requirements considered in the text are different. They are not conditioned on the presence of reasons.

comply, in this particular case, with these rational requirements, which promise justified beliefs and actions over the long run, if violating them, in this particular case, promises an additional justified belief or action?²⁹

A similar problem arises for the suggestion that one's reason for complying with rational requirements is that doing so is necessary to preserve oneself as a believer and agent. Take the insight shared by Davidson (2004, pp. 196–7) and Korsgaard: that complying with rational requirements is constitutive of being a believer and agent, so that to violate them is, at the limit, to cease to be a believer or agent at all.³⁰ Using this insight in a way that neither Davidson, nor Korsgaard would likely endorse, a nonreductionist might argue that one's reason to comply with rational requirements is to preserve oneself as a believer or agent. The problem is that it rarely threatens one's survival as a believer or agent to violate a rational requirement in any particular case. Yet we don't check how close one is to the point of no return before invoking the 'ought' of rationality, and we don't think that the 'ought' of rationality weakens the further from that point one is.

This problem would not arise, of course, if our reason for complying with rational requirements were an intrinsic reason, a reason to avoid irrationality for its own sake. That would be a reason to comply with

²⁹ Broome (2003, lec. 1, pp. 6–7) explores this view, but Broome (ms) decisively rejects it, for similar reasons. The only sense in which rationality is normative is that it is the case that we ought to have the rational faculty: a disposition to comply with rational requirements. It is not the case that we ought, in any particular case, to exercise this faculty. He agrees that failure, in a particular case, to exercise this faculty justifiably occasions criticism. But he offers a different explanation of the object of this criticism. 'If you do not satisfy some particular requirement of rationality, that is evidence that you do not have the rational faculty, at least not to the highest degree. So it is evidence that you are failing to achieve something you ought to achieve. It is therefore grounds for real criticism.' Two comments on this. First, it doesn't seem to do justice to the fact that such criticism is typically directed at what is going on here and now. Second, this criticism does not seem to be, at root, normative. It might be the case that I cannot make myself fully rational, in which case it would not be true that I ought to make myself fully rational. Nevertheless, my being irrational in a particular instance would justifiably occasion criticism. In this case, one might still say that it would be *better* if I were fully rational. I am defective in a certain way. (Compare intelligence. It would be better if I were more intelligent, in precisely the same way. I'd be more likely, over the long run, to arrive at the truth. But I do not have reason to make myself more intelligent, nor should I be criticized for failing to make myself more intelligent, since there's nothing I can do about it.) So it seems to me that this kind of criticism must be at root evaluative, rather than normative. I go on to consider, in section 4, the possibility that rationality is evaluative rather than normative.

³⁰ Jones (2003) and (ms b) explores a related possibility. Complying with rational requirements, like B+, B-, I+, and I-, that govern the relation between attitudes and beliefs about reasons for them, she suggests, may be constitutive of being a 'reason-responding agent.' Although Jones entertains the idea that being a reason-responding agent is something of intrinsic value, she, like Davidson and Korsgaard, does not appear to intend to use this idea to account for reasons to comply with rational requirements.

each and every rational requirement. But is there such a reason? We have intrinsic reasons to care about persons, relationships, justice, art, science, the natural environment, and so on, for their own sake. All of that is familiar enough. But is being subjectively rational another substantive value that we actually weigh against these others?

A nonreductionist might argue that it is, by referring once again to Davidson's and Korsgaard's insight, that complying with rational requirements is part of what it is to be a believer and agent. It is, so to speak, what sets us above the brutes. It is a valuable status. And I have conceded that its value gives us reason to preserve it. If violating a rational requirement would lead one to cease to be a believer or agent, then one has instrumental reason to not to violate it. The difficulty, as we have seen, is that this may not be true in any particular case. But—one might say—if this status is valuable, then we have reason not only to preserve it, but also to manifest it whenever the opportunity presents itself. This account, if plausible, would be of the right kind. It would give us a reason to comply with each and every rational requirement. However, it seems rather precious and unreal. When was the last time that a reason to express your status as a rational being weighed with you? I take it that this is why Davidson and Korsgaard do not employ in *this* way the idea that when we conform to rational requirements we realize our status as rational beings. They do not think that we deliberately do or ought to aim at expressing our rational status as a kind of substantive goal. Their point is instead that, at the limit, to fail to express our rational status is to fail to have that status. This point does not help the nonreductionist, as we have seen, because we are rarely, if ever, at the limit.

Korsgaard's view—as distinguished from Davidson's—is not simply that one is a believer and agent only if one *conforms* to the principles of rationality, but also that one is a believer and agent only if one *accepts* the principles of rationality, only if one is *committed* to them even in the cases in which one does *not* conform to them. Hence, it isn't really possible for a believer or agent to ask herself, with any sincere doubt, 'Ought I to follow the requirements of rationality?' To be a believer or agent is, so to speak, always already to have answered that question in the affirmative. Now I don't deny that *Korsgaard* can give this explanation of the normativity of rationality. What I deny is that a *nonreductionist* who seeks to explain the normativity of rationality *in terms of the normativity of reasons* can give it. The nonreductionist believes that an answer to the question, 'Why ought I to *X*?' must offer a substantive reason for *X*-ing, e.g., that *X*-ing would prevent suffering, or advance

the frontiers of knowledge. But Korsgaard doesn't aim to provide a *substantive reason* for following the principles of rationality. She aims instead to show that a commitment to following the principles of rationality is a *presupposition* of having the capacity to respond to substantive reasons in the first place. For the nonreductionist to put forward *this* explanation would be for him to concede that the normativity of rationality *cannot* be understood in terms of the normativity of reasons. Of course, the nonreductionist might take Korsgaard's claim out of its proper context and try to derive a substantive reason from it. One is an agent only if one is committed to following the principles of rationality: or, as the nonreductionist would likely put it, only if one *believes that one ought* to follow the principles of rationality. In so far as one has substantive reason to preserve oneself as an agent, the nonreductionist might argue, one has substantive reason to believe that one ought to follow the principles of rationality. But this would show only that one has (nonepistemic) reason to *believe* that one ought to follow the principles of rationality. It would not show that one ought to *follow* the principles of rationality, which is what the nonreductionist is trying to show.³¹

I do think, to anticipate, that rationality may be a kind of executive virtue, like courage. Every display of courage is, in some sense, admirable, even that of a mafioso in his gambit to take control of the local drug trade. Might this mean that there is at least one reason for the mafioso to make the gambit: namely, that it would be an admirable display of courage? Likewise, it might be said, every display of rationality is, in some sense, admirable, even that of someone who is mistaken about what he ought to do. Might this mean that there is at least one reason for him to do what he believes he ought to do: namely, that it would be an admirable display of rationality? I am inclined to say 'no.' To do something just because it would be courageous seems fetishistic. Of course, in special circumstances, there might be instrumental reasons to display courage. It might be the best way of cultivating the virtue in oneself, or showing others that one has it. But it does not seem true, as a general rule, that the mere fact that an act would display courage is itself a reason to do it. We don't imagine the mafioso's long-suffering brother, the priest, advising him: 'Well, there's at least this to be said for

³¹ Here I should record my doubt that it is necessary for being a believer or agent that one accepts that one ought to follow the requirements of rationality, or, for that matter, that one has any attitude toward them. As will become clear in the next section, we do not typically, if ever, reflect on the requirements of rationality when we comply with them. I would agree, however, that it is partly constitutive of being a believer or agent that, when in fact the requirements of rationality require one to have an attitude, one accepts that one ought to have that attitude. The Transparency Account explains why this is.

doing it: in so far as you do it, you will not be being deterred by fear.' Likewise, to do something just because it would be subjectively rational seems fetishistic. It is not true, as a general rule, that the mere fact that an act would display subjective rationality is itself a reason to do it.

I can't claim to have examined every proposal about what reasons we might have to comply with rational requirements, but I have tried to examine those that seem initially most promising. The inadequacy of some of these proposals, and the implausibility of others, is at least some further ground, in addition to the bootstrapping problem, to doubt that we have reasons to comply with rational requirements.

3. Third argument that there are no reasons to be rational: we typically do not, and in some cases cannot, reason from the fact that rationality requires compliance to compliance

A final ground for doubt that the fact that rationality requires certain attitudes is, or entails, a reason to have those attitudes is that we do not, and many cases cannot, reason from that consideration to compliance: to forming those attitudes.

It seems clear enough that we *can* comply with rational requirements without reasoning from them. First, we typically *do* comply with rational requirements without doing so on the grounds that rationality requires us to. When a person satisfies the antecedent of B+, for example, he believes that there is conclusive evidence that *p*. If he then goes on to form the belief that *p*, thereby complying with B+, he does so on the grounds of the evidence he believes there is, not on the grounds of his recognition that, given that he believes that there is conclusive evidence, it would be irrational of him not to believe that *p*.³² This second reason would seem superfluous from his point of view. Given that he believes that there is conclusive evidence that *p*, he believes that he already has all of the reason he needs. What could the thought that rationality requires it add? Moreover, to deny that we can comply with rational requirements without reasoning from them—that is, to assert that we can comply with rational requirements only by reasoning from them—would risk the kind of regress famously illustrated by Carroll's (1895) dialogue between the tortoise and Achilles.³³

³² A point stressed by Scanlon (2003, p. 20) and (ms).

³³ I say 'risks,' because the regress requires not only the claim that (1) one complies with a rational requirement, *R*, only if one reasons correctly from the content of the belief that *R* rationally requires compliance, to compliance, but also the claim that (2) one reasons correctly from the content of attitude *A* to attitude *B* only if one thereby complies with a rational requirement, *R*, governing *A* and *B*. In order for the reasoning in (1) to be correct, according to (2), one must comply

We can, and typically do, comply with rational requirements, then, without reasoning from them. Thus, if the fact that rationality requires compliance is a reason to comply, it is a reason that we can, and typically do, ignore. In this respect, it would be an odd sort of reason. But there is an even more serious problem. There are certain rational requirements that we *cannot* reason from. Thus, if the fact that rationality requires compliance is a reason to comply, it is a reason that we not only *do* not reason from, but also *cannot* reason from. This raises doubt about whether it can be a reason at all.³⁴

Searle (2001, p.104) puts the crux of this doubt with characteristic clarity and directness: ‘you have to be able to reason with reasons.’ Somewhat less clearly and less directly, I will express it as the ‘Generalized Internalism Requirement’, or

GIR: That *p* is a reason for *A* to *R* only if it is possible for *A* to *reason* from the content of the recognition that *p* is a reason for *A* to *R*³⁵ to *R*-ing.³⁶

Many philosophers, some of whom are otherwise opposed to Williams’s (1981) ban on ‘external reasons,’ have taken something like the GIR (at least as it applies to reasons for action), to be the ban’s kernel of truth. Parfit (1997, p. 114, n. 28) quoting Williams (1981, p. 102), takes it to be ‘uncontroversial’ that ‘if certain facts are claimed to provide nor-

with some further rational requirement *R'*, and in order to comply with *R'*, according to (1), one must reason correctly from the content of the belief that *R'* rationally requires compliance. This last piece of reasoning will need its own rational requirement, *R''*, and so on.

³⁴ Shah (ms) presents a similar, but far more developed, argument. In addition to laying out the central argument in greater detail, Shah offers deeper justification for and explanation of its premises. He also presents an independent line of argument for evidentialism, rooted in the claim of Shah (2003) that truth is the ‘standard of correctness’ for belief.

³⁵ Some might want to leave out the phrase, ‘is a reason for *A* to *R*.’ It does not matter for the present argument. Just as one cannot reason from the content of the belief that rationality’s requiring one to comply is a reason for one to comply to compliance, so too one cannot reason from the fact that rationality requires one to comply to compliance.

³⁶ One might deny, against Aristotle and Searle (2001), that reasoning can conclude in an action. The closest that it can come is what Searle (1983) calls an ‘intention-in-action.’ But an action comprises an intention-in-action and the corresponding bodily movement, and the relation between an intention-in-action and the bodily movement is not a further stage of reasoning. It is not a rational relation. It is purely causal. The GIR would then imply that there are no reasons for action, strictly speaking, only reasons for intentions-in-action. I don’t think this is a ground for rejecting the GIR. If one insists on speaking so strictly that reasoning cannot conclude in an action, that one cannot rationally respond to a reason by acting, that there is no rational relation between recognizing a reason and acting, then one ought, in consistency, to speak so strictly that there are no reasons for action. The considerations favouring the first restriction seem equally to favour the second restriction.

mative reasons, it must be true that '*people sometimes act for those reasons.*' Wallace (1999, pp. 217–8) takes it to be equally uncontroversial that 'if agent *A* has reason *r* to perform action *X*, and *A* is properly aware that *r* obtains, then *A* must be motivated to do *x*, on pain of irrationality,' which entails the GIR (at least as it applies to reasons for action). The GIR does justice, I think, to Williams's two basic insights about normative claims of the form that *A* has reason to *R*: that they have a potential explanatory dimension, and that they say something different from evaluative claims of the form that it would be good, in some way, if *A R*-ed.³⁷ In my view, the internalism requirement that Williams invoked to honour these insights was unnecessarily strong. A person might have a reason to *X* even if *X*-ing does not serve one of his present intrinsic desires. It is enough that the reason is the sort of consideration on the basis of which a person in his situation might reason to *X*-ing. And Williams's insights are not special to reasons for action. They apply to reasons in general.

Consider the claim that the fact that flat feet will keep me from being drafted is a reason for me to have flat feet. It is difficult to read this claim literally. I don't have reason to have flat feet as I might have reason to protest the draft, to believe that the war cannot be won, or to desire that it come to an end. The claim seems at best elliptical: I have reason *to want* to have flat feet, or *to act* so as to make it the case that I have flat feet. Why is this? What makes having flat feet different from acting, or believing, or desiring? One wants to answer, 'Because I can realize that having flat feet would keep me from the draft and that that would be a good thing, but it is not up to me to have flat feet on that basis.' The contrasting sense in which it is 'up to me' to act, believe, or desire, however, cannot be that of my having voluntary control over it. I cannot decide to believe or desire something any more than I can decide to have flat feet. On the other hand, it is not enough to say that belief is 'up to me' in the sense that my recognition of reasons for a belief can cause, more or less immediately, my believing it. For, in principle, my recognition of my 'reason' to have flat feet could cause, more or less immediately, the appropriate changes in my feet. The sense in which acting, believing, and desiring are 'up to me' is that I can come to act, believe, and desire on the basis of recognizing reasons for acting, believing, and desiring, in such a way that I thereby count as reasoning. By contrast, if my recognition of my 'reason' to have flat feet caused the

³⁷ Williams (1995, pp. 39–40) thus complains that the external reasons theorist cannot say what 'the difference [is] supposed to be between saying that the agent has a reason to act [in some way], and saying... that it would be better if [he so] acted'.

appropriate changes in my feet, that would be a welcome event, but it would not be reasoning.

If this is correct, then there are no ‘state-given’ reasons for belief, only ‘object-given’ reasons for belief, to use terminology introduced by Derek Parfit. An object-given reason for an attitude is supposed to be a consideration that recommends the attitude by indicating something about its object. For example, an object-given reason for the belief that *p* is a consideration that recommends that belief by indicating that it is true that *p*. And an object-given reason for the intention to *X* is a consideration that recommends that intention by indicating that *X*-ing is a worthwhile thing to do. A state-given reason for an attitude, by contrast, is supposed to be a consideration that shows that having the attitude itself would be good in some way. For example, a state-given reason for intending to drink a toxin might be that an eccentric billionaire will give me a million dollars if I intend to drink it—whether or not I actually drink it. And a state-given reason for believing that I am good looking might be that I will be more confident if I believe it—which does nothing, alas, to show that it is true. Now, arguably, it is possible to reason to an *intention* from the recognition of a state-given reason for it. But it is not possible to reason to a *belief* from the recognition of a state-given reason for it: that is, a reason that one does not take to be evidence that it is true. This is not reasoning.³⁸ This is not to deny that people sometimes do form beliefs as a causal consequence of entertaining state-given reasons. There are cases of self-deception and wishful thinking. But these beliefs are not related to the recognition of those state-given reasons as links in a single chain of reasoning.³⁹ If one cannot reason from state-given reasons for belief, then the GIR implies that there are no state-given reasons for belief, properly so called. State-given ‘reasons’ for belief are like ‘reasons’ for flat feet. At most, they are

³⁸ Broome (2003b, lecs 2 and 3) appears to agree with this. For his purposes there, however, the more important point is that, as a contingent fact, we cannot form beliefs and intentions on the basis of beliefs that we have reason to have them, not that, if we could, we would not be reasoning correctly. Moreover, neither point is meant to be limited to beliefs that one has state-given reasons. For example, Broome believes that one cannot intend to *X* on the basis of a belief that one has an object-given reason for that intention. One can intend to *X* only on the basis of a belief that one has reason to *X*. I am not convinced that it is not in our power to intend to *X* on the basis of state-given reasons, let alone on the basis of object-given reasons, at least so long as we believe that it is otherwise possible for us to *X*. And it is not clear to me that this could not count as reasoning. If I am wrong about this, however, then that only strengthens the argument of this section.

³⁹ It can be proper reasoning to form, on the basis of considering state-given reasons for a belief, an intention to act in a way that will lead one to form that belief. Here, however, the process of reasoning culminates not in a belief, but instead in an intention to act in a certain way. The chain of events that follows the intention—the process by which one’s actions result in one’s belief—is not itself part of the same reasoning process.

reasons to want to believe that p , or reasons to act so that I come to believe that p .⁴⁰

Why does this show that there are no reasons in general to comply with rational requirements? Because if there were reasons to comply with rational requirements, these reasons would have to be state given. A subject would form an intention to X not because X -ing seemed to him a worthwhile thing to do, but instead because having that intention was necessary to avoid irrationality. Similarly, a subject would form a belief not because it seemed to him, in view of the apparent evidence, likely to be true, but instead because having that belief was necessary to avoid irrationality. Since there are no state-given reasons for belief, it follows that, at least where compliance involves belief, the fact that rationality requires compliance is not a reason for one to comply. This is a third ground for doubting that there are reasons to conform to rational requirements, in addition to the bootstrapping problem and the obscurity about what these reasons might be.

4. Why rational requirements are not simply evaluative

If the normativity of rationality is not that of a reason, then what might it be? At this point, one might say that it has been a mistake all along to suppose that rational requirements are, strictly speaking, normative—or, as one might say, deontic, or response guiding. Indeed, it is a mistake to suppose that the principles of rationality are, strictly speaking, requirements or demands addressed to the subject to whom they apply. They are instead standards of assessment or evaluation. The primary role of principles of rationality is not in first-person deliberation, but in third-person appraisal. You, as evaluator, use them to see how I measure up to some standard. But I, as subject, do not follow them in forming my responses.

What kind of evaluation might be at issue? One possibility, which Tim Scanlon has raised, is that standards of rationality define a certain kind of proper functioning. In failing to conform to them, one manifests a certain kind of functional defect.⁴¹ The evaluative notion at the core of claims about rationality, on this view, is that of how something

⁴⁰ Here one might protest: 'It makes more sense to say I ought to believe that p because it would be good for me to, than to say I ought to have flat feet because it would be good for me to.' I grant that it *sounds* better. But I suspect that this is only because whereas it is never correct to say that I ought to have flat feet, it is sometimes correct to say that I ought to believe that p : namely, when I have object-given reasons to believe it.

⁴¹ Scanlon (2003) mentions this possibility, but Scanlon (ms) goes on to express doubts about it.

is supposed to work. More specifically, rationality might be understood as the proper functioning of a particular kind of system for producing beliefs and intentions. If the subject believes that he has conclusive reason for some belief, for example, but does not form it, then the system is not performing as it 'ought' to. This does not mean that the subject has *reason* to form the belief. But it does mean that his system of belief formation is *malfunctioning*.

One worry with this proposal is that irrationality is a failing of the person, not of some subpersonal system. This raises a doubt about whether a person's being irrational can really be equivalent to a system's malfunctioning. Another, more serious problem is that it is unclear what it means to say that the system is 'supposed' to function in this way. In cases of intention and natural selection, it seems plausible to reduce the 'supposed' of a function to something descriptive. To say that an artifact is supposed to function in this way is to say, roughly, that its designer or user intends for it to function in this way. To say that an animal's organ is supposed to function in this way is to say, roughly, that that organ was selected for because, in so far as it functioned this way, it contributed to the fitness of the animal's ancestors. Yet the 'supposed' of rationality cannot be explained in either of these ways. It cannot be explained in terms of intention, because we were not intentionally designed, and we are subject to standards of rationality whether or not we ever intended to be. And it cannot be explained in terms of natural selection, because we are subject to standards of rationality whether or not conforming to those standards promoted our ancestors' fitness or promotes ours.⁴² It is an a priori truth, if you like, that we are, as believers and intenders, subject to standards of rationality. It is not something that evolutionary biology could confirm or disconfirm.

Perhaps there is a way of honouring this point. It is *constitutive* of belief and intention, it might be said, that each is the product of a system with a characteristic function. I won't take issue here with the idea that what it is for something to be belief, say, is for it to be the product of a system with a constitutive function. But even so, a worry arises. The most plausible candidate for the constitutive function is truth. This would suggest, in turn, that the system is functioning properly when it is responsive to evidence: that is, when it conforms to the norms of

⁴² Contrast, however, Foot (2001), who argues, drawing on Thompson (1995), that there is a notion of proper biological functioning that cannot be reduced in terms of selective advantage. She proposes to understand the normativity of specifically human rationality in terms of such a function, a possibility which I do not discuss here.

objective epistemic rationality.⁴³ The difficulty is that these norms of *objective* epistemic rationality are not the norms of *subjective* epistemic rationality that we are trying to account for. Appeal to a constitutive function of belief might explain the sense in which belief is 'supposed' to be responsive to *evidence*. But how then is it to explain the sense in which belief is 'supposed' to be responsive to (potentially false) *beliefs* about evidence? The mystery about the relation between 'ought' of reasons and the 'ought' of rationality seems only to have been replaced by a mystery about the relation between the 'supposed' of reasons and the 'supposed' of rationality.

It might be suggested that the difference between the 'supposed' of rationality and the 'supposed' of reasons is a difference between part and whole. If someone forms a belief in response to the belief that the evidence supports it, when in fact it does not, then part of the system is functioning as it is supposed to, although the system as a whole is not. Evaluations of subjective rationality would then be evaluations not of the whole system, but instead of certain parts: such as a part whose function is to form beliefs in response to beliefs about evidence. It remains to be explained, first, how the constitutive function of the system as a whole determines the functions of its parts in general and, second, how it determines the function of one such part in particular: that of forming beliefs in response to beliefs about evidence. But perhaps these things can be explained.

Indeed, I think that an account along something like these lines may be correct, although it seems more natural to construe it not in terms of the proper functioning of a system, but instead in terms of a virtue—a substantively good way for a person to be.⁴⁴ Some virtues are dispositions to recognize certain kinds of reasons and to respond accordingly. Kindness is a disposition to respond to the needs of others, for example, and justice is a disposition to respond to considerations of fairness. As we have seen, however, rationality cannot be a virtue of this kind. It is not a disposition to act on a special class of reasons. However, rationality might be understood as a kind of executive virtue. Executive virtues, like courage and tenacity, are not dispositions to recognize and respond to a special kind of reason. They are, instead, dispositions that help one to execute one's beliefs about one's reasons, whatever they might be, or to execute one's intentions, whether or not one believes there are reasons for them. Courage, for example, is a disposition not to be deterred

⁴³ As Burge (2003) suggests.

⁴⁴ Compare Dancy (2000, pp. 60–70); and Svarvarsdottir (2003).

by fear from doing what one believes one ought to do, or what one, perhaps akratically, intends. Suppose—as my Transparency Account will suppose—that rationality consists in having the attitudes that one believes that one has reason to have. Then rationality seems a kind of executive virtue. It is a disposition to execute one's beliefs about one's reasons for and against one's attitudes.

We evaluate positively the executive virtues and their manifestations, independently of what they are used for. We admire the courage that the mafioso displays in lunging for power, even though we don't believe that it's something that he ought to do. Likewise—the suggestion would go—we admire the rationality that someone displays in intending to *X*, because he believes that he ought to *X*, even though we do not believe that in fact he ought to *X*. It is an interesting question why we do this. For present purposes, however, it may be enough that we do. What we want to know is in what sense we evaluate positively rationality and its manifestations. The suggested reply is: In the same sense in which we evaluate positively the executive virtues and their manifestations.

I find this evaluative account plausible, as far as it goes. Often, when we claim that people are irrational, in general or on specific occasions, we may be simply evaluating them, for lacking an executive virtue or for failing to manifest it. But this evaluative account cannot be the whole story. It leaves out a normative dimension, which is manifest both 'from the outside,' in advice to the effect that it would be irrational of the subject to fail to believe or intend something, and 'from the inside,' in his experience of being bound by a rational requirement to believe or intend it.

Consider, first, this normative dimension from the outside. It is not always true that when we claim that someone is irrational, we are only appraising him. Sometimes we seem to be saying something normative, or deontic, or response guiding: something in the register of advice, rather than assessment. An atheist might say to a racist believer: 'Look: I think you're nuts to believe in God, let alone that He created anyone. But given that you believe that God created all people equal, and given that you agree that people whose skin is a different colour from yours are people, you *ought* to believe that He created them equal too. It would be irrational of you not to.' Claims of this kind do not seem to be grading the addressee. Indeed, in advance of seeing how the addressee will respond, the addresser may not yet have grounds for grading him at all. Instead, such claims seem to be recommending to him, in some way, a certain response.

Consider, now, this normative dimension from the inside. Often when we ourselves are subject to rational requirements, we feel that we ought to respond as they require, or at least we can be brought to feel this when the right things are called to our attention. Recall, or imagine, for example, what clear-eyed akrasia is like. In feeling this normative ‘pressure,’ we are not, to be sure, reflecting on the fact that rationality requires us to have attitude *A*, much less judging that this is itself reason to have *A*. But we do feel that we ought to have *A*. In other words, we do not think that we ought to give the response that rationality requires of us *as such*, much less that we ought to give it *because* rationality requires it of us. Rather, *of* the response that rationality requires of us, we think that we ought to give it. All the same, this thought is something distinctly normative, rather than evaluative. We are not stepping back from ourselves and praising ourselves for rationality, or criticizing ourselves for irrationality. For one thing, this would be unnaturally reflexive. For another—to repeat a point just made—in advance of seeing how we will respond, we have no grounds for grading ourselves in the first place.

In sum, there is a kind of ‘ought’ that comes naturally to us when subjective rationality is at stake, and this ‘ought’ seems, like the ‘ought’ of having conclusive reason, to express something normative rather than evaluative. To this extent, Broome was absolutely right. This brings us to the heart of the puzzle. On the one hand, we need to explain how rationality *is* normative, in the sense that claims about rationality can function as advice, not just as appraisal, that rational requirements can seem to guide responses, not simply to rate them. On the other hand, we need to explain how rationality is *not* normative, in the sense that there are reasons to be rational. In what sense ‘ought’ we to comply with standards of rationality, if not in the sense that we have reason to comply with them?

Might nonreductionists bite the bullet and accept that alongside the primitive ‘ought’ of reasons, there is a distinct, primitive ‘ought’ of rationality? On a certain line of thought, this might not seem troubling: ‘There are lots of “oughts”’, one might say. ‘There is the “ought”’, for example, of the rules of chess. One “ought”, according to the rules of chess, not move a pawn backwards. Clearly, this does not mean that one has conclusive reason not to move a pawn backwards. Why can’t we simply say that there is a special “ought” of rationality, which likewise does not imply conclusive reasons? The problem is that the ‘ought’ of chess is purely classificatory. To say that someone ‘ought’, in this sense, to conform to the rules of chess is simply to say that in so far as he is to

do what will count as playing chess, he must not move his pawn backwards. When we say that someone ought to conform to the standards of rationality, however, we are not making the merely descriptive claim that in so far as he is to count as rational, he must conform to the standards that define what counts as rational. We are recommending, in some way or other, that he comply with those standards. This is not to deny that we sometimes recommend that people conform to purely classificatory standards, like the rules of chess. If we have gone to great lengths to get Fisher to Reykjavik to see him beat Spassky, then we might advise him to do what will count as playing chess. Since he ought not disappoint us, he ought not move his pieces willy-nilly. But this cannot be the sense in which one 'ought' to comply with standards of rationality. To begin with, Fisher's 'ought' obtains only because of his particular situation. In another situation, it might not be the case that he ought not move his pawn backwards. (On his own time, if he found it a rewarding thing to do, he might play 'house' with the pieces, with the pawn a child returning home from school.) The 'ought' of rationality, by contrast, applies whatever one's particular situation. More fundamentally, this 'ought' reflects a substantive reason that Fisher has to do what will count as playing chess: namely, not to disappoint us. As we have seen, there is no substantive reason to be rational. In sum, the concession that there is a second, primitive normativity of rationality is more significant than the analogy to the classificatory 'ought' of chess might suggest.

It is a concession, I suspect, that normative nonreductionists are not eager to make. This is not simply for reasons of parsimony, but also because of the difficulty of understanding how we are to be governed by these two autonomous 'ought's. As we have seen, what one is rationally required to believe or intend will sometimes conflict with what one in fact has conclusive reason to believe or intend. If reasons and rational requirements give rise to autonomous, primitive 'ought's, then what *really* ought one to do in such conflicts? Some compromise between what one has reason to do and what one is rationally required to do? Or is it rather that we cannot even ask what one *really* ought to do, only what one ought-in-the-reasons-sense to do and what one ought-in-the-rationality-sense to do? These are unpalatable alternatives, and nonreductionists might have hoped to avoid a forced choose between them.⁴⁵ But is there any way for them to avoid it?

⁴⁵ I am indebted to Nadeem Hussain for urging me to articulate my misgivings about this 'double-primitive' option.

5. Explaining the normativity of rationality: The Transparency Account

Nonreductionists may be able to avoid it, at least if—and I grant that this is a big ‘if’—rational requirements take a certain form. Consider what I will call the ‘core requirements’:

C+: If one believes that one has conclusive reason to have *A*, then one is rationally required to have *A*.

C–: If one believes that one lacks sufficient reason to have *A*, then one is rationally required not to have *A*.

Now suppose that all rational requirements are ultimately derived from these core ones. Then a possibility emerges for explaining the normative dimension of rationality. Let us begin by considering it ‘from the outside.’ When we tell someone, in the register of advice, rather than that of appraisal, that he ought rationally to have attitude *A*, or that it would be irrational of him not to have it, suppose that we are simply pointing out that he satisfies the antecedent of C+. We are making the descriptive, psychological claim that he believes that he has conclusive reason for the attitude. We are telling him, as we might put it, that *from his point of view, or as it seems to him*, he has conclusive reason to have the attitude. Thus, when we tell him that he ‘ought rationally’ to have attitude *A*, we are not ourselves offering him a reason to have *A*. How, then, are we advising him to have *A*? By drawing his attention to a reason for *A* that *he believes* he has.⁴⁶ Thus, a second-person charge of irrationality, ‘But you ought to believe it! It would be irrational of you not to!’ says, in effect: ‘Look: from *your* point of view, you have reason to believe it!’ Likewise, a third-person charge of irrationality, ‘He ought to believe that *p*. It would be irrational of him not to,’ says, in effect, ‘From *his* point of view, he ought to believe that *p*. That is something that someone could point out to him, if given the chance.’⁴⁷

⁴⁶ Michael Smith notes that one can sincerely advise *S* that *S* ought to have *A* only if one believes that *S* can have *A*. My view, however, cannot explain this. For it is perfectly coherent for me to point out to *S* that *S* believes that *S* ought to have *A* even if *I* do not believe that *S* can have *A*. This is, I think, simply a difference between ‘objective’ and ‘subjective’ advice. When I advise *S* ‘objectively,’ committing myself to the claim that *S* ought to have *A*, then I must believe that *S* can have *A*. But when I advise *S* ‘subjectively,’ intending only to draw *S*’s attention to what *S* believes, namely that *S* ought to have *A*, then I need only believe that *S* believes that *S* can have *A*. I am not committed to the claim that *S* ought to have *A*, so I am not committed to the claim that *S* can have *A*.

⁴⁷ Pam Hieronymi and Seana Shiffrin have independently suggested to me an ingenious generalization of this. Instead of saying, ‘If one believes that one has conclusive reason to have *A*, then one is rationally required to have *A*,’ I should say, ‘If what one believes, if true, would provide con-

Why does the advisee experience this as advice? More generally, why does being subject to a rational requirement such as C+ feel normative ‘from the inside’? Because a reason that someone believes he has is, from his point of view, a reason *simpliciter*. In other words, given what the antecedent of C+ is, it will always seem to someone to whom C+ applies that he has reason to comply with it. This reason is not that given that he satisfies the antecedent of C+, it will be irrational of him not to have A, and irrationality is something to avoid. The reason is instead the reason for A that, in virtue of satisfying the antecedent of C+, he already believes he has. From the first-person standpoint, the ‘ought’ of rationality is transparent. It looks just like the ‘ought’ of reasons. It is only from the second- or third-person standpoint that the ‘ought’ of rationality and the ‘ought’ of reasons come apart. For it is only from a standpoint other than the subject’s that it is possible to distinguish what attitudes he has reason to have from what attitudes, as it seems to him, he has reason to have.

This account does not appeal to a second primitive: an additional normative concept beyond that of a reason. The (seeming) normative force of the ‘ought’ of rationality derives from a (seeming) reason, the reason that the subject believes he has.

Even though the normative concept is that of a reason, however, this account does not entail that there are reasons to be rational. On this account, to say that someone ‘ought rationally’ to have some attitude is to say that, as it seems to him, he has reason to have it. And to say that, as it seems to him, he has reason to have it is not to say that, in fact, he has reason to have it. So this account does not lead to bootstrapping.

clusive reason to have A, then one is rationally required to have A.’ Accordingly, when we say, ‘You ought to A; it would be irrational of you not to,’ we mean not (in general), ‘As it seems to you, you ought to A,’ but instead, ‘As it seems to you, something is so, and (although you may not have realized it) if that is so, then you ought to A.’ On the original account, we can preface rational advice by informing the advisee of the normative consequences of what he believes and leading him to conclude, on that basis, that he ought to have A. But, according to that account, it makes sense to deliver that advice—to say, ‘You ought to A; it would be irrational of you not to’—only if the advisee comes to agree that his beliefs have those normative consequences and thus concludes that he ought to A. According to the proposal that Hieronymi and Shiffrin offer, such agreement is not necessary. This proposal shares the principal virtues that I claim for the Transparency Account; it doesn’t appeal to a second primitive, it doesn’t entail that there are reasons to be rational, it doesn’t raise the mystery of what such reasons might be, and it explains why we do not reflect on rational requirements when we comply with them. While I would be otherwise happy to adopt the proposal, I have one misgiving. I don’t believe that people are being irrational in failing to have A simply because they do not know that what they believe, if true, provides them with conclusive reason to have A. Ignorant perhaps, but not irrational.

Nor does it raise the mystery of what reasons we have to be rational. The apparent reasons to which we draw the addressee's attention are not reasons to be rational. They are just the garden-variety reasons—prudential, moral, epistemic, and so on—that he believes he has to do *X* or to believe that *p*.

The Transparency Account also explains why we do not reflect on rational requirements, let alone reason from them, when we comply with them. Take a subject who satisfies the antecedent of C+: that is, who believes that she has conclusive reason, *R*, for attitude *A*. Suppose that we say to her: 'You ought rationally, given what you believe, to have attitude *A*.' On what was once Broome's view, we are informing her of another reason to have *A*: namely, that, given what she believes, it would be irrational of her not to have *A*, and she must not be irrational. Having been informed of this, she now sees herself as having two reasons for attitude *A*: *R* and the reason that it would be irrational of her not to have *A*. This is an odd state of mind. To begin with, the second reason—that she is rationally required—is superfluous from her point of view. If she believes that she has conclusive reason for *A*, then, from her perspective, she already has conclusive reason to have it. What could the thought that rationality requires it add? Moreover, if *A* is a belief then it is not a consideration that she can reason from. She recognizes it, but it is, for her, inert. It can have no rational influence on how she forms or revises her attitudes. Hence, if the GIR is correct, then it is not a reason at all. On the Transparency Account, by contrast, we are not informing her of any additional reason to have *A*, when we say to her: 'You ought rationally, given what you believe, to have attitude *A*.' We are simply drawing her attention to *R*, the reason that she already believes she has.

The Transparency Account, however, explains only the normative force—or, better, apparent normative force—of the core requirements. So it explains the (apparent) normativity of rationality as a whole only if every instance of irrationality is such because it violates a core requirement. This is, in effect, Scanlon's (1998, pp. 25–30) view: that we take subjective rationality, as a whole, to consist in having attitudes that cohere with one's beliefs about the reasons for them. But offhand it is not at all clear that this view is defensible.⁴⁸ For example, the requirement to intend the apparent means if one intends the end, does not seem to have the form of either of the core requirements. It does not seem to involve a belief that one has or lacks reasons for an attitude.

⁴⁸ The view seems to be shared by Davidson (2004). Scanlon (ms) develops this view in a way that answers many of the worries that I go on to discuss in the text.

Indeed, it does not seem to involve any beliefs at all. Likewise, rational requirements to have logically consistent beliefs do not seem reducible to the core requirements. After all, one is rationally required to hold logically consistent beliefs, no matter what their subject matter: that is, whether or not they are beliefs about reasons for attitudes. Indeed, it is not even clear that the requirements that I used to argue against Broome's former account, namely I+, I-, B+, and B-, are captured by C+ and C-. I+ involves the belief that one has conclusive reason to *do* something, not that one has conclusive reason to have the attitude of *intending* to do it. And B+ involves the belief that *there is* conclusive evidence that something *is* the case, not that *one has* conclusive reason to *believe* that it is. Nevertheless, I think that, despite appearances, these rational requirements can either be derived from the core requirements, or be shown to be defeasible approximations to them. So I think that the Transparency Account can explain the (apparent) normative force of rationality in general. But this is a task for a different paper.⁴⁹

If the Transparency Account is correct, however, then normative nonreductionists can claim to have achieved the reverse of what Smith and Korsgaard aspire to. Whereas they have sought to explain reasons in terms of rationality, nonreductionists can claim to have succeeded in explaining rationality in terms of reasons. However, nonreductionists have succeeded in this, if they have, only by borrowing one of the central insights of Korsgaard's constructivism. This is that the normativity of rationality has its source in the first-person perspective of deliberation. Rational requirements get a 'grip on us' only from within the outlook framed by the questions, 'What ought I to believe? What ought I to do?'⁵⁰

313 Emerson Hall
Harvard University
Cambridge, MA 02138
USA
kolodny@fas.harvard.edu

NIKO KOLODNY

⁴⁹ At a few points in this paper, I have offered some thoughts on how this might be done. In subsection 1.3, I argued that I+ and I- are approximations to C+ and C-, and in the last paragraphs of subsection 1.9, I suggested how one might handle requirements of logical consistency. Scanlon (ms) argues that many familiar rational requirements flow from the core requirements. In particular, he offers a broadening of the view that subjective rationality is a matter of having attitudes that cohere with one's 'attitude-directed' judgements about the reasons for and against those attitudes. Subjective rationality, he suggests, is more often a matter of having attitudes that cohere with one's 'content-directed' judgements, like the judgements about evidence and reasons for action that serve the antecedents of B+, B-, I+, and I-.

References

- Arpaly, N. 2000: 'On Acting Rationally against One's Best Judgment'. *Ethics*, 110, pp. 488–513.
- Bermudez, J. and A. Millar (eds) 2003: *Reason and Nature: Essays in the Theory of Rationality*. Oxford: Oxford University Press.
- Broome, J. 1999: 'Normative Requirements'. *Ratio*, 12, pp. 398–419.
- 2001: 'Normative Practical Reasoning'. *Proceedings of the Aristotelian Society*, suppl. vol. 75, pp. 175–93.
- 2003a: 'Practical Reasoning' in Bermudez and Millar (eds), 2003, pp. 85–112.
- 2003b: *Reasoning*. Brown-Blackwell Lectures.
- 2004: 'Reasons' in Wallace, Pettit, Scheffler, and Smith (eds) (2004) pp. 28–55.
- ms: 'Rationality'.
- Burge, T. 2003: 'Perceptual Entitlement'. *Philosophy and Phenomenological Research*, 67, pp. 503–48.
- Carroll, L. 1895: 'What the Tortoise Said to Achilles'. *Mind*, 4, pp. 278–80.
- Dancy, J. 2000: *Practical Reality*. Oxford: Oxford University Press.
- Davidson, D. 2004: 'Incoherence and Irrationality' in *Problems of Rationality*. Oxford: Oxford University Press, pp. 189–98.
- Foot, P. 2001: *Natural Goodness*. Oxford: Oxford University Press.
- Harman, G. 1986: *Change in View*. Cambridge, MA: MIT Press.
- Hatzimoysis, A. (ed.) 2003: *Philosophy and the Emotions*. Cambridge: Cambridge University Press,
- Hursthouse, R., G. Lawrence, and W. Quinn (eds) 1995: *Virtues and Reasons*. Oxford: Oxford University Press.
- Jones, K. 2003: 'Emotion, Weakness of Will, and the Normative Conception of Agency' in Hatzimoysis (ed.) (2003), pp. 181–200.
- ms a: 'Intellectual Self-Trust'.

⁵⁰I owe a special debt to Tim Scanlon, for several reasons, not the least of which is that this paper arose from our discussion of a draft of his 'Structural Irrationality.' I am also greatly indebted to John Broome, Ruth Chang, Jonathan Dancy, Luca Ferrero, Barbara Herman, Nadeem Hussain, Pam Hieronymi, Karen Jones, Doug Lavin, John MacFarlane, Derek Parfit, Glenn Ross, Sam Scheffler, Seana Shiffrin, Michael Smith, Sigrun Svarvarsdottir, Sergio Tennenbaum, Jay Wallace, Wai-hung Wong, and two anonymous reviewers for *Mind*, for comments on drafts, many of them quite extensive and detailed. Earlier versions of this paper were presented over the past year in the Moral Psychology Conference at Franklin and Marshall College, the Reasons and Rationality Conference at the Research School of Social Sciences of the Australian National University, and in colloquia at Ohio State University, University of California at Los Angeles, and University of California, Berkeley. I would like to thank participants at these events, too many to attempt to list by name, for very helpful criticism. Finally, I would like to express my gratitude to the Australian Research Council, which supported much of the writing of this paper through Grant DP0343704.

- ms b: ‘What Should We Want from a Theory of (Human) Rationality?’
- Kamm, F. 2000: ‘The Doctrine of Triple Effect and Why a Rational Agent Need Not Intend the Means to his End’. *Proceedings of the Aristotelian Society*, suppl. vol. 74, pp. 21–39.
- Kavka, G. 1983: ‘The Toxin Puzzle’. *Analysis*, 43, pp. 33–6.
- Korsgaard, C. 1996: *The Sources of Normativity*. Cambridge: Cambridge University Press.
- Lavin, D. 2004: ‘Practical Reason and the Possibility of Error’. *Ethics*, 114, pp. 424–57.
- Parfit, D. 1997: ‘Reasons and Motivation’. *Proceedings of the Aristotelian Society*, suppl. vol. 71, pp. 99–130.
- ms: *Rediscovering Reasons*.
- Pollock, J. L. and J. Cruz 1999: *Contemporary Theories of Knowledge*, 2nd edn. Lanham, MD: Rowman and Littlefield.
- Raz, J. 1999: *Engaging Reason*. Oxford: Oxford University Press.
- Scanlon, T. M. 1998: *What We Owe to Each Other*. Cambridge, MA: Harvard University Press.
- 2003: ‘Metaphysics and Morals’. *Proceedings and Addresses of the American Philosophical Association*, 77, pp. 7–22.
- ms: ‘Structural Irrationality’ in G. Brennan, R. Goodin, and M. Smith (eds), a volume of essays in honour of Philip Pettit.
- Searle, J. 1983: *Intentionality*. Cambridge: Cambridge University Press.
- 2001: *Rationality in Action*. Cambridge, MA: MIT Press.
- Shah, N. 2003: ‘How Truth Governs Belief’. *Philosophical Review*, 112, pp. 447–82.
- ms: ‘A New Argument for Evidentialism’.
- Smith, M. 1994: *The Moral Problem*. Oxford: Blackwell.
- 2004: ‘The Structure of Orthonomy’. *Philosophy*, suppl. 55, pp. 195–216.
- Svarvarsdottir, S. 2003: ‘Evaluations of Rationality’. *Southern Journal of Philosophy*, 41, pp. 121–36.
- Thompson, M. 1995: ‘The Representation of Life’ in Hursthouse, Lawrence and Quinn (eds) 1995, pp. 247–96.
- Wallace, R. J. 1999: ‘Three Conceptions of Rational Agency’. *Ethical Theory and Moral Practice*, 2, pp. 217–42.
- Wallace, R. J., P. Pettit, S. Scheffler, and M. Smith (eds) 2004: *Reasons and Value: Themes from the Moral Philosophy of Joseph Raz*. Oxford: Oxford University Press, pp. 28–55
- Williams, B. 1981: ‘Internal and External Reasons’ in *Moral Luck*. Cambridge: Cambridge University Press, pp. 101–13.

——1995: 'Internal Reasons and the Obscurity of Blame' in *Making Sense of Humanity*. Cambridge: Cambridge University Press, pp 35–46.