

The Myth of Practical Consistency

Niko Kolodny

It is often said that there is a special class of norms, 'rational requirements', that demand that our attitudes be related to one another in certain ways, whatever else may be the case.¹ In recent work, a special class of these rational requirements has attracted particular attention: what I will call 'requirements of formal coherence as such', which require just that our attitudes be formally coherent.² For example, we are rationally required, if we believe something, to believe what it entails. And we are rationally required, if we intend an end, to intend what we take to be necessary means to it. The intuitive idea is that formally incoherent attitudes give rise to a certain normative tension, or exert a kind of rational pressure on each another, and this tension, or pressure, is relieved just when one of the attitudes is revised. As John Broome observes, these requirements are, by their nature, 'wide scope', which is to say that there is no particular attitude that one must have or lack in order to satisfy them. This is because they require just formal coherence, and there is no particular attitude that one must have or lack in order to be formally coherent.

For many, these requirements of formal coherence are the only sure thing in the domain of normativity. Either, as Humeans say, there is nothing else. Or, as Kantians say, whatever else there is must be a construct, or relative, of requirements of formal coherence.³ Yet, in his much-discussed paper, 'The Myth of Instrumental Rationality', Joseph Raz puts this view on the defensive.⁴ He denies, in particular, that there is a rational requirement of formal coherence that governs means and ends, such as:

Means-End (ME): One is rationally required (if one intends at t to E , believes at t that one will E only if one intends at t to M , then to intend at t to M).

Raz offers two observations to explain the initial plausibility of *ME*. First, he makes an observation about the structure of reason:

Facilitative Principle: If there is *pro tanto* reason for one to E , then there is at least as much *pro tanto* reason for one to take some sufficient means to E .⁵

Second, Raz observes that rationality inherits the structure of reason. We are rational, Raz suggests, insofar as we exercise the capacities that enable us, in general, to respond correctly to reason (even if these capacities occasionally lead us astray). If these capacities are to enable us, in general, to respond correctly to reason, they must involve some sensitivity to the patterns of reason that the

Facilitative Principle describes. Accordingly, when one's rational capacities register (perhaps incorrectly) that there is reason for an end and so lead one to intend that end, those capacities will usually also register (perhaps incorrectly) that there is reason to take what appear to be necessary means to the end and so lead one to intend those apparent means. No wonder, then, that we usually judge someone to be responding rationally when he intends the apparent necessary means to his ends. But this has nothing to do with *ME*: with a special principle that demands means-end coherence as such. That is a myth. There is only the structure of reason and our sensitivity, insofar as we are rational, to that structure.

At certain points, Raz seems to suggest that the myth of instrumental rationality is a special case of a broader myth, that there is, normatively, a set of principles that requires, or, descriptively, a special faculty that somehow strives for, formal coherence as such.⁶ My aim, in this paper, is to explore this suggestion, by applying Raz's approach to a different alleged requirement of formal coherence as such:

Simple Intention Consistency (C): One is rationally required (either not to believe at *t* (that if one *X*'s, then one does not *Y*),⁷ or not to intend at *t* to *X*, or not to intend at *t* to *Y*).⁸

C should be familiar from other philosophical connections, beyond the theory of rationality. In the philosophy of mind, it is often said that *C* distinguishes intentions from desires, and in the philosophy of action, that *C* distinguishes acting with an intention from acting intentionally.⁹ For Kant, *C* is pivotal in applying the categorical imperative. It is the sense in which one cannot both will an immoral maxim and will that it become a universal law.¹⁰ And *C* may be crucial to versions of noncognitivism that interpret apparent normative beliefs as intentions. If *C* is false, then intentions may not 'rule out' inconsistent intentions in the way in which apparent normative beliefs rule out inconsistent apparent normative beliefs.¹¹

I start by asking whether we can explain the intuitions that seem to recommend *C* along Razian lines: in terms of the structure of reason and our sensitivity, insofar as we are rational, to that structure. After some preliminary remarks about reason, in section 1, I present, in section 2, a Cognitivist Version of the Error Theory for *C*. In light of some doubts about the thesis of Confidence, which the Cognitivist Version requires, I present, in section 3, a Noncognitivist Version. In large part, the Noncognitivist Version spells out an analytical framework for the positive account of reason for intention in Hugh McCann's article, 'Settled Objectives and Rational Constraints'. As McCann observes, and as I discuss in sections 4 and 5, there are cases in which inconsistent intentions violate no requirement of reason or rationality.¹² These 'Exceptions' confront *C* with what I call the Problems of Extension and Conflict. In section 6, I discuss a third problem, the Problem of Normativity, which is independent of the Exceptions. In section 7, I complete the Error Theory, by applying Raz's

observation that rationality inherits the structure of reason. In section 8, I conclude by discussing the upshot for a more descriptive, or constitutive, telling of the myth: that what it is for a state to be an intention is, in part, for it to coordinate action by means of dispositions that strive for consistency, and other kinds of formal coherence, as such.

1. Preliminary Remarks About Reason

In large part, *C* seems plausible because it would explain:

Violation intuition (V): If one believes at *t* (that if one *X*'s, then one does not *Y*), intends at *t* to *X*, and intends at *t* to *Y*, then one violates some requirement.

The core of the Error Theory for *C* is that the requirement in *V* is not the wide-scope rational requirement of formal coherence *C*, which is satisfied by revising any of the three attitudes, but instead some narrow-scope requirement of reason, which is satisfied only by revising certain of the attitudes. In other words, it is not the tension among *the three attitudes* that requires a revision, but instead the tension between certain of those attitudes and *the reason against them* that requires a revision.

By the mass noun, 'reason', I mean not the faculty, but instead the collection of considerations on the basis of which we settle the questions, 'What to believe? What to choose?'¹³ Reason, in other words, is the material of first-person deliberation, within which these questions are asked. So understood, one's reason for a belief or intention¹⁴ might be a function of one's attitudes (such as beliefs, desires, nondoxastic perceptual states), or of facts independent of one's attitudes. While I favor the latter possibility, I leave it open for the purposes of this paper. It depends on the substantive question of what matters within deliberation.¹⁵

The explanation of *V* might seem straightforward. We first observe:

Certain Failure: If one intends to *X*, intends to *Y*, and believes that if one *X*'s, then one does not *Y*, then either one's intention to *X* does not succeed, or one's intention to *Y* does not succeed, or one's belief that if one *X*'s, then one does not *Y* is false.

And then we claim (i) if one's belief that *p* is not true, then one violates a requirement of reason in having that belief, and (ii) if one does not succeed in one's intention to *X* at *t*, then one violates a requirement of reason by intending, at any time, to *X* at *t*. The difficulty is that one need not violate a requirement of reason by having a false belief, or by not succeeding in an intention. This is because the questions to be resolved within deliberation are whether one's belief is *likely, given the evidence*, to be true or false and whether one's intention is *likely, given the evidence*, to bring about something worthwhile.¹⁶ If one violates a

requirement of reason in believing that p , it is because p is not sufficiently likely, in light of the evidence, to be true, relative to the values of acquiring truth and avoiding falsity, and if one violates a requirement in intending to X , it is because intending to X does not make it sufficiently more likely, in light of the evidence, that one will X , relative to the cost of the intention and the value of X -ing.

These claims about when one violates a requirement of reason are true at least if we assume, as I will, that there are no 'state-given' reasons for or against belief, and no 'state-given' reasons for intention. That is:

Evidentialism: The only reason to believe p is that the evidence makes it likely that p , and the only reason not to believe p is that the evidence makes it likely that p is false,

and

Aimism: There is reason for one to intend at t to X only insofar as (i) intending at t to X raises, causally,¹⁷ the probability that one X 's and (ii) there is reason for one to X .

Thus, when A and B are patterns of intentions that differ only in that one intends to X in A , but does not intend to X in B , there is more reason to have A than to have B , if there is, *only because* having A raises the probability that one X 's and there is reason for one to X . Aimism is compatible with there being state-given reason *not* to intend to X : reason not to intend to X other than that it increases the probability that one X 's, which there is reason not to do. Intending X can have other costs, besides raising the probability that one X 's. For one thing, intending X may make one more likely to take costly means to X .¹⁸ For another, intending X may lower one's chances of succeeding in other intentions, whether or not it succeeds in getting one to X (because, e.g., it distracts one). These state-given reasons *against* intending seem ubiquitous and uncontroversial (which may help to explain why they are scarcely discussed in the literature).¹⁹

We often need to think in terms of reason not for or against individual intentions, but instead for or against *patterns* of intentions. For example, I may have sufficient reason (to intend to book a room and to intend to attend the conference), but not sufficient reason (to intend to book a room and not to intend to attend the conference) (because this is likely to lead me to forfeit a deposit on a room I never use), and not sufficient reason (not to intend to book a room and to intend to attend the conference) (because attending the conference is likely to be fruitless, if I am distracted and irritable from nights of tossing and turning on park benches). Quite often, the question, 'Should I intend to X ?' will have no determinate answer. The real question is whether I should intend to X in addition to, or in the absence of, intending some Y .

Why is this? In general, one's reason to intend to X depends on the likelihood that other conditions, that affect one's reason to X , or whether one will succeed in X -ing, will or will not obtain. I have more reason to intend to picnic tomorrow, for

example, if it is less likely to rain. These conditions are sometimes the actions of a person. I have more reason to intend to attend this conference if you are more likely to attend. And sometimes the person is the agent himself. I have more reason to intend to find accommodation where the conference is being held, if I am more likely to attend the conference. And I am more likely to attend if I intend to. This is the phenomenon that T. M. Scanlon, in another context, calls the 'predictive significance' of intention.²⁰ When I intend something, this changes what the future is likely to hold. And, in general, changes in what the future is likely to hold, whether or not those changes are brought about by my intentions, change what I have reason to do and/or to intend. In sum, my intending *X* can affect whether I have reason to intend *Y*. Thus, it makes no more sense to ask whether I should intend to *Y* in abstraction from whether I intend to *X* than it makes sense to ask whether I should put out the washing to dry in abstraction from whether it will rain. However, unlike the rain, my intending or not intending to *X* is itself something for which I can have or lack reason. So in deciding what intentions I ought to have, I cannot take it as given that I do intend to *X*, as I take it as given that it will or will not rain. The upshot of both points is that we often cannot simply consider my reason to intend to *X*, but must consider my reason (to intend to *X* and to intend to *Y*) and my reason (to intend to *X* and not to intend to *Y*).²¹

Scanlon's idea has far-reaching consequences. It explains, I think, many of the phenomena that requirements of formal coherence as such on intention are invoked to explain.²² The idea is not, as with requirements of formal coherence as such, that we are required if we intend *X*, then to intend (or not to intend) things that we believe are formally related to *X*. The idea is instead that if we intend *X*, then this often changes what the future is likely to bring, which in turn changes what else we have reason to intend. The appeal to predictive significance differs in several ways from the appeal to requirements of formal coherence as such. First, what matters for requirements of formal coherence is what we *believe* about the relations between the intentions, whereas what matters for predictive significance is how they are *likely* to affect one another.²³ Second, what matters for requirements of formal coherence is intending *X as such*, whereas what matters for predictive significance is what intending to *X means for the future*. In some cases, intending *X* (e.g. intending to become King of France in 2008, or intending *X* when there is compelling evidence that I do not intend *X*) may *not* affect the likely future. In other cases, conditions other than intending to *X* (e.g. my self-trust that I will *X* even if I don't intend it now, or my knowledge that I will be tortured or coerced into *X*-ing) may affect the likely future in the *same* way. Third, when having one intention affects my reason for another intention, it does so by instantiating an *attitude-independent* condition of a normative principle: a condition that can be satisfied by factors that are not my own attitudes. For example, there is the normative principle:

Effectiveness: If some condition makes it less likely that my intention to *X* will bring about my *X*-ing, then I have less reason to intend to *X*.

This condition need not be that I have some attitude. It might be that my hamstring is pulled, that the gas tank is empty, that the electorate is hostile, that the bank has frozen the account, and so on. On occasion, however, this fact may be the fact that I also intend to *Y*. When this is so, I can have weaker reason (to intend to *X* and to intend to *Y*) than (to intend to *X* and not to intend to *Y*). When one intention puts me under a requirement of formal coherence to revise another, by contrast, it does so by instantiating a condition that only an intention can satisfy. Finally, requirements of formal coherence are wide-scope. That is, one never needs to have or lack any particular attitude in order to satisfy them. By contrast, one usually needs to have or lack a particular attitude in order to conform to one's reason to have a certain pattern of intentions. If one has conclusive reason, for example, (to intend *X* and not to intend *Y*), then one can comply with this requirement of reason only by intending *X* and not intending *Y*.²⁴

2. The Cognitivist Error Theory for *C*

With these preliminary remarks out of the way, we can ask how the requirement in *V* might be some narrow-scope requirement of reason. Suppose we assume:

Confidence: If one intends to *X*, then one believes that one will *X*.²⁵

Now consider two facts about the structure of reason:

Transmission: When *q* is a logical consequence of *p*, then the evidence that *q* is at least as strong as the evidence that *p*,

which one might think of as a theoretical analogue of Raz's Facilitative Principle, and,

Stronger Evidence: One has sufficient reason (reason permits one) to believe *p* only if the evidence that *p* is stronger than the evidence that not-*p*.²⁶

Together, these facts make it plausible²⁷ that:

Comment about Reason Patterns for Belief: Either one lacks sufficient reason (reason requires one not) to believe *p*, or one lacks sufficient reason to believe *q*, or one lacks sufficient reason to believe (if *p*, then not-*q*).

Together with Confidence, this entails:

Cognitivist Comment about Reason Patterns (RC): If one believes (that if one *X*'s, then one does not *Y*), intends to *X*, and intends to *Y*, then either one believes without sufficient reason that one will *X*, or one believes without sufficient reason that one will *Y*, or one believes without sufficient reason that (if one *X*'s, then one does not *Y*).

RC says, in other words, that if one has inconsistent intentions, then one has some belief that violates a narrow-scope requirement of reason. Thus, *RC* provides an

explanation of *V*: the intuition that, if one has inconsistent intentions, then one violates some requirement. According to the *Cognitivist Version of the Error Theory*, *RC*, and not *C*, explains *V*.²⁸

The weakness of the Cognitivist Version is that it relies on Confidence. It might be argued, on the contrary, that one can intend something while not believing that one will do it. As Bratman illustrates, one may intend to carry out a chancy, last-ditch rescue operation without much confidence that one will succeed. Or one might intend to stop at the bookstore on one's way home while knowing that one has a tendency to go on 'automatic pilot', which may well lead one to forget.²⁹

The advocate of Confidence may reply that these counterexamples are only apparent. First, he may claim that the agent does not intend *to X*, but instead intends *to try to X*.³⁰ But what is it *to try to X*? It is not simply taking all of the steps that one would take if one were to execute an intention to *X*. Suppose that, innocent of boomerangs, Captain Cook follows Amaro's step-by-step instructions, surprised that it returns to knock the tricorn off his head. Cook does not try to knock the tricorn off, one wants to say, because he does not take those steps with the goal of knocking it off, because he is not guided, in his activity, by that aim. What else can this missing state of mind be, but an intention to hit it?³¹ If trying to *X* consists, in part, in intending to *X*, and if the agent in the rescue case not only intends to try, but actually tries, to rescue, then he does in fact intend to rescue, despite doubting that he will succeed.

The second proposal is that the agent does not *intend to X*, but instead 'aims', or 'weakly intends' to *X*.³² To explain away the apparent counterexamples, however, this 'aiming' must have every property the intention to *X* has except those properties that are ruled out by the absence of an expectation of success. That is, it must be the case that (*modulo* any differences stemming from the absence of such an expectation):

- (i) aiming to *X* and intending to *X* are similar enough to be supported to the same extent by the same reasons;
- (ii) at least in cases in which one believes that reason underdetermines choice, one can aim at will, in the way in which one can intend, but not desire, at will;
- (iii) present-directed aiming to *X*, like present-directed intending to *X*, guides the agent's conduct toward the end of *X*-ing;
- (iv) aiming to *X* has the same role in instrumental reasoning as intending to *X*;
- (v) future-directed aiming to *X* has the same tendency to stability as future-directed intending to *X*;
- (vi) the belief that one will not *X* causes the revision of aiming to *X*, as it causes the revision of intending to *X*; and
- (vii) self-knowledge that one aims to *X* is like self-knowledge that one intends to *X*.

These similarities encourage the following suspicion. If we cut the mind at its joints, we find only 'aimings'. Often, as when one has faith in one's abilities and the world's cooperation, one's aiming is accompanied by the expectation that one will succeed. But, other times, when one lacks such faith, one's aiming is not.³³ Of course, we might reserve the word, 'intention', for aimings of the former sort. But then Confidence would be a merely linguistic thesis. Whether or not it would find support in ordinary usage,³⁴ it would seem of limited philosophical interest.

3. The Non-Cognitivist Error Theory for C

While I harbour this suspicion about Confidence, my position in this paper is agnostic and disjunctive. If Confidence is true, then the Cognitivist Version of the Error Theory is available. If Confidence is false, then a Non-cognitivist Version is available. In place of *RC*, this second error theory appeals to:

RN: Usually, either reason requires one not to believe (that if one *X*'s, then one does not *Y*), or reason requires one not (to intend *X* and to intend *Y*) —that is, usually, either one lacks sufficient reason to believe (that if one *X*'s, then one does not *Y*), or one has more reason (to intend to *X* and not to intend to *Y*) than (to intend to *X* and to intend to *Y*), or one has more reason (not to intend to *X* and to intend to *Y*) than (to intend to *X* and to intend to *Y*), or one has more reason (not to intend to *X* and not to intend to *Y*) than (to intend to *X* and to intend to *Y*).

However, the explanation of *RN* is rather more complicated than the explanation of *RC*.

When it is the case that if one *X*'s, then one does not *Y*, *why* is it the case? Most commonly, one is in a:

Decision-requiring Situation: One must eventually decide between *X* and *Y* in order to achieve either.

One must decide to drive or walk, to see the comedy or the thriller, to have Chinese or Italian, to buy whole or skim, to take the job at City College or State U., to vacation in the mountains or by the sea, to vote Democratic or Republican, in order to do any of these things. Decision-requiring Situations are necessarily *Intention-preventing Situations*, which is to say that something (i) that would be necessary for fulfilling one's intention to *X* and (ii) that would be brought about by that intention (such as deciding in favor of *X*-ing, or successfully *X*-ing) would prevent one from fulfilling one's intention to *Y*.

However, there are also *Decisionless Situations*, in which one need not decide between *X* and *Y*, in order to achieve either.³⁵ One can 'let the world decide', as Bratman puts it.³⁶ For example, one can embark on the ongoing, multi-stage process of applying to North State, intending to gain admission there, and do the

same with respect to South State, until either school accepts one. Among these Decisionless Situations, we can distinguish between:

Decisionless, Intention-preventing Situations: One can let the world decide. Some factor

- (i) without which one will not fulfill one's intention to *X* and
- (ii) that, if it obtains, is brought about by one's intention to *X*

prevents one, if it obtains, from fulfilling one's intention to *Y*.

and

Decisionless, Non-intention-preventing Situations: One can let the world decide. Some factor

- (i) without which one will not fulfill one's intention to *X* and
- (ii) that, if it obtains, is *not* brought about by one's intention to *X*

prevents one, if it obtains, from fulfilling one's intention to *Y*.

Suppose that the admissions process is organized so that *once* one is admitted to one State campus, one's application is withdrawn from the other. Then the situation would be Intention-preventing, because the factor that, if it obtains, prevents one from being admitted to (say) North State—namely, being admitted to South State—will be, if it obtains, *brought about by the intention* to be admitted to South State. By contrast, suppose that one is *already* ineligible for admission at one campus, but one has forgotten which campus. Then the situation would be Non-intention-preventing, because the factor that, if it obtains, prevents one from being admitted to North State—namely, being ineligible for admission to North State—will *not* be, if it obtains, brought about by the intention to be admitted to South State.

RN is certainly true if either there is not sufficient reason to intend to *Y*, whether or not one intends to *X* (or vice-versa), or there is not sufficient reason to intend additionally to *Y*,³⁷ whether or not there is sufficient reason to believe that if one *X*'s, then one does not *Y* (or vice-versa), or there is not sufficient reason to believe that if one *X*'s, then one does not *Y*. But such cases teach us nothing about why one should lack sufficient reason for inconsistent intentions in particular. The interesting question is why *RN* is true in cases in which one has sufficient reason (not to intend to *X* and to intend to *Y*) (or vice-versa), but lacks sufficient reason (to intend to *X* and to intend to *Y*), because there is sufficient reason to believe that if one *X*'s, then one does not *Y*.

Suppose, then, that there is sufficient reason to believe that if one *X*'s, then one does not *Y*. Suppose, as a notational convenience, that there is no less reason to *X* than to *Y*. And suppose, for the time being, that intending additionally to *Y*

makes it more likely that one *Y*'s than would not intending additionally to *Y*. If this is an Intention-preventing Situation, then intending additionally to *Y* in turn makes it less likely that one *X*'s, because it increases the probability of that which would prevent one from *X*-ing. This is an instance of the more general phenomenon of:

Competition: additionally intending to *Y* makes it less likely that one *X*'s.

Avoiding Competition is often what is meant by talk of 'coordinating' one's intentions. Notice that intentions can compete without being jointly unrealizable,³⁸ and that intentions can be jointly unrealizable without competing.³⁹ Importantly, the belief involved in *C* is that the two intentions are *not jointly realizable*, not necessarily that they are *in Competition*. Now suppose that Competition is *Total*: the increase in the probability that one *Y*'s is matched by at least as great a decrease in the probability that one *X*'s. In Decision-requiring Situations, it seems plausible to assume that Competition is Total. Intending to see the comedy tonight, in addition to intending to see the thriller tonight, for example, does not make it more likely that one will (either see the comedy tonight, or see the thriller tonight). If there is:

Stronger Reason: more reason to *X* than to *Y*,

then there is not sufficient reason to intend additionally to *Y*. Intending additionally to *Y* makes it more likely that one achieves something that one has less reason to do, only by making it less likely, by the same degree, that one achieves something that one has more reason to do. Even when Competition is Less than Total, as it may be in some Decisionless, Intention-preventing Situations, Stronger Reason may make it the case that one lacks sufficient reason to intend additionally to *Y*.

Suppose, however, that there is:

Equal Reason: the reasons for *Y*-ing are just as strong as those for *X*-ing.

Then, even if Competition is Total, this means at most that there is no reason to intend additionally to *Y*. Unless there is some *pro tanto* reason *not* to intend additionally to *Y*, there might be sufficient reason to intend additionally to *Y*. This is also the case if, contrary to our earlier supposition, additionally intending to *Y* does not make it more likely that one *Y*'s.

Often, *pro tanto* reason *not* to intend additionally to *Y* is provided by the fact that it may lead one to take *Extra Means*. For example, additionally intending to see the thriller may lead one to buy an extra ticket in advance. This only adds to the cost of the evening, without increasing its benefits.

In Decision-requiring situations, further *pro tanto* reason *not* to intend additionally to *Y* is provided by the mechanism of *Decision Undoing*. This is important, because, in some situations, there may not be Stronger Reason or Extra Means. For example, if one cannot buy tickets in advance, then the means to seeing the comedy might be, until the moment of truth, the same as the means

to the thriller: scheduling a babysitter, taking an early dinner, driving to the multiplex, and so on. So why does one have reason, *before* the moment of truth, *not* to intend additionally to see the thriller? The suggestion is this. Deciding between the comedy and the thriller is, so to speak, a bit of work that one must *eventually* do. Of course, there may be no reason to do this work sooner rather than later. One can simply intend (either to see the comedy, or to see the thriller) and leave one's decision until the moment of truth. But if one *does* form the intention to see the comedy, full stop, before the moment of truth, then one thereby does the work of deciding between them. At this point, to intend additionally to see the thriller would only undo this work. One would be back where one started: having to decide, eventually, between the comedy and the thriller. This is some reason *not* to intend additionally to see the thriller. Of course, it is not *much* reason not to intend it. But *any* reason is enough, it would seem, given that there is no reason *to* intend it.

But what if, against reason, one *does* intend additionally to see the thriller? So far, this account identifies no reason against *keeping* those inconsistent intentions. Keeping them does not somehow *further* undo one's decision; that is water under the bridge. Keeping them only postpones making a new decision between them. As we noted, there may be no reason not to postpone one's decision. Still, there may seem something amiss in keeping those inconsistent intentions.

If something is amiss here, it has nothing in particular to do with inconsistency. The same problem arises in consistent cases. Suppose that reason requires me to see the comedy, because my editor needs a review of that and only that movie. At the moment, I intend only (either to see the comedy, or to see the thriller). As before, there is no reason for me not to postpone, until I reach the box office, forming an intention to see the comedy. However, if I pause now to form some particular intention, then it is a kind of Decision Undoing to decide to see the thriller. I waste the opportunity, forcing myself to revisit the matter later on, in order to revise my intention to see the thriller and to form an intention to see the comedy. Suppose that I make this mistake, forming an intention to see the thriller, while not forming an intention to see the comedy. Just as before, there seems something amiss in *keeping* the intention to see the thriller. Yet, unlike before, this intention is not inconsistent with any other intention I have. Perhaps we should say that, in general, there is something amiss in *keeping* a pattern of intention that reason required one *never* to adopt, even if there is nothing amiss in *postponing* the adoption of one of the alternative patterns of intention that reason requires one *eventually* to adopt. Or perhaps we should give up our sense that anything is amiss. Either way, appealing to C is not the right response to the phenomenon.⁴⁰

To sum up, Decision-requiring Situations corroborate RN, because they involve Total Competition, which gives one no reason *to* intend additionally to Y, and Decision Undoing (and perhaps also Stronger Reason and Extra Means), which gives one reason *not* to intend additionally to Y. Usually, Decisionless, Intention-preventing Situations also corroborate RN, since usually such cases involve sufficient Competition and either sufficiently Stronger Reason to X, or sufficiently costly Extra Means in intending to Y, that, all things considered, one

lacks sufficient reason to intend additionally to *Y*. Notice that Competition, which is the principal mechanism behind *RN*, is an instance of the predictive significance of intention, in which one intention, by changing what is likely to come, changes one's reason for other intentions.

4. The Exceptions and the Problems of Extension and Conflict

However, this leaves out certain *Exceptions*. This is why *RN* is qualified by 'usually'. Consider Decisionless, Intention-preventing Situations, in which there is Less than Total Competition, Equal Reason, and no Extra Means. In Bratman's ingenious *Video Game*, one guides with each hand a different missile toward a different target, knowing that as soon as one hits one target, the game will shut down, preventing one from hitting the other.⁴¹ Intending to hit target *Y* makes more likely not only states in which one hits *Y* when, if one had not hit *Y*, one would have hit *X* (i.e. in which missile *X* is on target, but *Y* gets there first), but also states in which one hits *Y* when one would not have hit *X* anyway (i.e. in which missile *X* is off course). So the increase in the probability that one hits *Y* is greater than the decrease that intending to hit *Y* brings about in the probability that one hits *X*. Bratman's description suggests, and we can anyway assume, that there is Equal Reason for *Y* and that additionally intending to hit *Y* does not lead one to take Extra Means (because, say, the only cost is the time spent and that cost has already been sunk by one's intention to hit *X*).

In Decisionless, Non-intention-preventing Situations, there is even less reason why *RN* should be true. In the *Hidden Target*, for example, either target *X* or target *Y* is missing, but one does not know which, because something obscures each location. There is no Competition, since there are no states in which one hits *X* but would have hit *Y*, if one had not hit *X*.

It may be helpful here to compare the (Intention) Exceptions, in which one violates no requirement of reason in having inconsistent intentions, to Belief Exceptions, in which one violates no requirement of reason in having inconsistent beliefs. The classic, 'preface' example is of an author who knows that he has sufficient reason to believe each of several claims in a book, as well as to believe, since no book is without errors, the negation of their conjunction.⁴² In one respect, Belief Exceptions are easier to find than Intention Exceptions. Apart from exotic cases, there is nothing like Competition in the case of belief. *Having* a belief does not lower the probability that an inconsistent belief is true. In another respect, however, Belief Exceptions are harder to find than Intention Exceptions. Although the ratio of the value of acquiring a true belief on a question to the value of avoiding a false belief on that question may vary from question to question, it seems plausible that the ratio remains within certain limits. In particular, there is good reason to expect that the importance of acquiring true belief is never as great as the importance of avoiding false belief. Otherwise, it would not be true, for instance, that one should suspend belief when the evidence is equally balanced between *p* and not-*p*.⁴³ By contrast, the ratio of the

value of a fulfilled intention to the value of avoiding an unfulfilled intention need have no such upper bound. We can, by the proper selection of circumstances, make the value of succeeding in some intention—what it would achieve—arbitrarily great and the value of avoiding a failed intention—the cost of trying—arbitrarily small.

To illustrate, contrast the case of two inconsistent beliefs with the case of two jointly unrealizable intentions. Let:

S be the value of ‘succeeding’ in either attitude over not having it,
 F be the value of not having either attitude over ‘failing’ (or alternatively ‘not succeeding’) in it,
 $P(\text{max})$ be the probability of the attitude more likely to succeed succeeding, if one has only that attitude,
 $P(\text{min})$ be the probability of the attitude less likely to succeed succeeding, if one has only that attitude,
 $P(\text{both})$ be the probability of one attitude succeeding, if one has both attitudes.

If we assume that one has sufficient reason to have both attitudes only if that has at least as great an expected value as having either attitude on its own, then there would be sufficient reason for both attitudes only if:

$$P(\text{max})S - (1 - P(\text{max}))F \leq P(\text{both})(S - F) - (1 - P(\text{both}))2F,$$

or,

$$(1 - P(\text{both}) + P(\text{max})) / (P(\text{both}) - P(\text{max})) \leq S/F.$$

In the belief case, $P(\text{max}) + P(\text{min}) = P(\text{both})$, and, by definition, $P(\text{max}) \geq P(\text{min})$. So, $P(\text{both}) - P(\text{max}) \leq 1/2$, which in turn means that the left-hand side is ≥ 1 . If, as I suggested, $S/F < 1$ in the case of belief, then the inequality is never satisfied. In the intention case, by contrast, there is no reason to assume that $S/F < 1$. So we can satisfy the inequality by making S/F sufficiently large. This is why Belief Exceptions involving only two inconsistent beliefs are implausible, whereas Intention Exceptions involving only two jointly unrealizable intentions, such as Video Game, are possible.

As we have seen, the Exceptions require that RN be qualified. And these qualifications mean that RN cannot explain V . By the same token, the Exceptions suggest that V should be similarly qualified. Consider *Known* Exceptions, in which one *knows* that one has sufficient reason to believe that if one X 's, then one does not Y , and knows that one has sufficient reason for both intentions. In such cases, intuitively, one can have all of these attitudes without violating any requirement, whether of reason or of rationality.⁴⁴ What seems to be true is not V , therefore, but instead:

V' : Usually, if one believes that if one X 's, then one does not Y , intends to X , and intends to Y , then one violates some requirement.

And RN may explain V' .⁴⁵

According to the *Non-Cognitivist Version of the Error Theory*, we mistakenly appeal to *C* to explain *V'*, when in fact *RN* explains *V'*. This appeal to *C* is mistaken in two ways. As I will explain in section 6, it confuses a disjunction of narrow-scope requirements of reason with a wide-scope rational requirement of a disjunction of responses. But first, and most obviously, the appeal to *C* overlooks the Exceptions.

Overlooking the Exceptions leads to the *Problem of Extension*. *C* entails that it is *always* irrational to have inconsistent intentions. But sometimes, as in the Known Exceptions, it is not, intuitively, irrational.

Overlooking the Exceptions also leads to the *Problem of Conflict*. Consider *Strong Exceptions*, in which reason not only permits, but also requires, one to have inconsistent intentions. Suppose that, through no fault of one's own, the world confronts one with a Strong Exception. If *C* were correct, then rationality would require one *not* to have inconsistent intentions, whereas reason would require one *to* have inconsistent intentions. This would contradict the plausible principle of:

Ideal Compatibility: If one is required at *t* by reason to *X* and required at *t* by rationality not to *X*, then there should be some earlier time *t'* and response *R* such that if one had given at *t'* *R*, then one would not be both required at *t* by reason to *X* and required at *t* by rationality not to *X*.

Ideal Compatibility does not deny that there can be conflicts between reason and rationality. It allows that one may paint oneself into a corner: respond in such a way that going forward one cannot satisfy what rationality requires without violating what reason requires. What is harder to accept, and what Ideal Compatibility does deny, is that the *world* should paint one into a corner, or, to vary the metaphor, lay traps for one: that there should be conflicts between reason and rationality that one could not have avoided, no matter what one had done.

5. Replies to the Problems of Extension and Conflict

One might first try to respond to these problems by claiming that, in the Exceptions, *C* is overridden.⁴⁶ On one interpretation, this means that, in the Exceptions, *rationality* does not require one to satisfy *C*: that is, that it is not *irrational* to violate *C*. But it is mysterious why this should be. On another interpretation, it might mean the following. Rationality always requires one to satisfy *C*, and in the Non-Exceptions, this gives one reason sufficient to require one to satisfy *C*. However, in the Exceptions the reason to violate *C* overrides this reason, so that all things considered reason permits one to violate *C*.

This is no response at all to the Problem of Conflict, since it concedes a conflict between reason and rationality. And it is not a persuasive response to the Problem of Extension. First, the intuition is that inconsistency violates no requirement in the Exceptions, not even a requirement of rationality. Second,

there is no reason why the reason to satisfy *C* should be overridden in 'Weak Exceptions', in which reason otherwise permits, but does not require, an inconsistent pattern. Video Game, as originally described, seems such a case; one may try to hit both targets, instead of only one, but one is not required to do so. (The game is just as much fun either way.) In such cases, the reason to satisfy *C* ought to break the tie, so that all things considered reason forbids one from having the inconsistent pattern. But this is not what happens.

Next, one might try to respond to these problems by appealing to a distinction that Bratman draws between intending and 'endeavoring'. Against the Problem of Extension, one might grant that we find it plausible that, in the Known Exceptions, one can have *attitudes* with contents that cannot be jointly realized without violating a requirement. But one might argue that we are implicitly assuming that the attitudes are *endeavorings*. So the Known Exceptions are not counterexamples to *C*.

Even granting some distinction between intending and endeavoring, however, I find this unpersuasive. After all, we can *explicitly* consider an agent who has *intentions*—that is, states of the same kind as those that are uncontroversially called 'intentions'—in a Known Exception. And when we do, the verdict, to my mind, remains unchanged. He does not violate any requirement.

Against the Problem of Conflict, one might argue that in the Strong Exceptions, reason does not require one (to intend to *X* and intend to *Y*), but instead requires one (either to intend to *X* and to intend to *Y*, or to endeavor to *X* and to endeavor to *Y*, or to intend *X* and to endeavor *Y*, or to endeavor *X* and to intend *Y*). So one can satisfy both this requirement of reason and *C*.

This is more persuasive. If endeavoring is similar enough to intending to be supported to the same extent by the same reasons, then (apart, perhaps, from an exotic situation in which one can only intend, but not endeavor) reason never requires one (to intend to *X* and to intend to *Y*).

In any event, the distinction between intending and endeavoring is itself elusive.⁴⁷ One cannot say that the difference consists *only* in the *normative* difference that *C* applies to intentions, but not to endeavorings. This would not only beg the question, but also deny the supervenience of the normative on the nonnormative. What, then, might the nonnormative difference be? The replies to the Problems of Extension and Conflict require that endeavoring is almost identical, nonnormatively, to intending. For example, it must be the case that (*modulo* any differences stemming from (x) below):

- (i) endeavoring to *X* and intending to *X* are nonnormatively similar enough to be supported to the same extent by the same reasons;
- (ii) at least in cases in which one believes that reason underdetermines choice, one can endeavor at will, in the way in which one can intend, but not desire, at will;
- (iii) present-directed endeavoring to *X*, like present-directed intending to *X*, guides the agent's conduct toward the end of *X*-ing;

- (iv) endeavoring to *X* has the same role in instrumental reasoning as intending to *X*;
- (v) future-directed endeavoring to *X* has the same tendency to stability as future-directed intending to *X*;
- (vi) the belief that one will not *X* causes the revision of endeavoring to *X*, as it causes the revision of intending to *X*;
- (vii) self-knowledge that one endeavors to *X* is like self-knowledge that one intends to *X*; and
- (viii) relative to the same background beliefs, endeavoring to *X* and intending to *X* support to the same extent the expectation that one will *X*, on which further planning may be based.

We cannot plausibly claim that:

- (ix) Intentions, but not endeavorings, are such that, normally, when they are inconsistent, the belief that one is in a Non-exception causes their revision,

because, in light of (i), the belief that one is in a Non-exception will normally cause the revision of endeavorings as well. However, we might claim that:

- (x) Intentions, but not endeavorings, are such that, normally, when they are inconsistent, the belief that one is in an Exception causes their revision.

If the difference is (x), however, then it is not clear why we should believe that we have 'intentions', so understood. Normally, when we believe that we are in Non-Exceptions, our states are such that this belief causes their revision, and normally, when we believe that we are in Exceptions, our states are such that this belief does not cause their revision. Of course, this evidence is *consistent* with the hypothesis that we have both endeavorings and 'intentions', so understood, but normally avoid such 'intentions' when we believe we are in Exceptions. But it is *equally* consistent with the hypothesis that we only endeavor—that what Bratman proposes to call 'endeavorings' are simply what we ordinarily call 'intentions'.

Occam's Razor seems enough to break the tie in favor of this latter hypothesis. Moreover, there are two pieces of positive evidence against the former hypothesis. The first is simply that, when we decide what to do, we do not see ourselves as facing a further choice whether to bring this about by endeavoring or by intending.

The other piece of evidence is that, as Bratman observes, in Video Game, for some *Z*, one intends to *Z* target 1 and intends to *Z* target 2. Yet there seems nothing for '*Z*' to be, but 'hit'. Bratman rightly rejects the suggestion (1) that one simply intends (either to hit that target, or to hit the other).⁴⁸ Whatever it is that one intends with regard to each target, one believes that a *necessary* means to it is

an attempt to hit *that* target. But my trying to hit *that* target is not a *necessary* means to (either hitting this target, or the other).

Next, it might be said (2) that the agent intends only to *try* to hit each target. But, as we noted in our discussion of Confidence, trying to hit each target seems to involve intending to hit each target. It might be replied that trying involves not intending, but only endeavoring. But this is implausible. The present-directed intention to try (the intention to try that guides one's trying as one tries) would be a present-directed intention to be guided by a present-directed endeavoring to *X*. It seems implausible that the agent has this redundant, second layer of intention-like state.

Nor, finally, is (3) one's intention to hit target 1, *if one can*. This must be understood as an intention that *succeeds* if, because one has hit target 2, one *cannot* hit target 1. Otherwise the two intentions could not both succeed, and they would be inconsistent.⁴⁹ But does one succeed in one's intention (to hit target 1, if one can), if, because one has hit target 2, one cannot hit target 1? Not if an intention succeeds only if it *makes* its content true. For it is the *other* intention—the intention directed at target 2—that makes the content of this intention true, by making it the case that one cannot hit target 1. Might we say instead that an intention succeeds if its content is true, whether or not it makes its content true? Then one might simply wait until one hits target 2, and then form and effortlessly fulfill, after the fact as it were, the intention (to hit target 1, if one can). Whatever one intends in Video Game, it is not this.

6. The Problem of Normativity

While the last problem with *C*, the *Problem of Normativity*, is intensified by overlooking the Exceptions, it does not depend on it. For the Problem of Normativity, the second mistake in the appeal to *C* suffices: the mistake of confusing a disjunction of narrow-scope requirements of reason with a wide-scope rational requirement of a disjunction of responses.

RN is a disjunctive observation about actual, or possible, narrow-scope requirements of reason. *RN* observes that, in any Non-Exception in which one has inconsistent intentions, there is *at least one specific* revision of those attitudes that reason requires. One of the following must be the case, it says: *either* reason requires one (to intend to *X* and not to intend to *Y*), *or* reason requires one (not to intend to *X* and to intend to *Y*), *or* reason requires one (not to intend to *X* and not to intend to *Y*),⁵⁰ *or* reason requires one not to believe that if one *X*'s, then one does not *Y*. It is not the tension among *those attitudes* that requires a revision, but instead the tension between some of those attitudes and *the reason against them* that requires a revision. One satisfies this requirement of reason *only* by revising the specific attitudes that it requires one to revise, not by revising other attitudes in the inconsistent set.

By contrast, *C* is a wide-scope rational requirement of a disjunction of responses. One is rationally required, *C* says, to bring about the following

disjunctive state of affairs: that either one revises one's intention to *X*, or one revises one's intention to *Y*, or revises one's belief that if one *X*'s, then one does not *Y*. As a requirement of formal coherence as such, *C* demands just formal coherence, and formal coherence is achieved equally well by any of these revisions. Here, it is precisely the tension among those attitudes that requires a revision. And one satisfies *C*—one relieves that tension—by revising *any* attitude in the inconsistent set.

This raises a question: What there is to be said, to or within the first-person standpoint of deliberation, for satisfying *C*: for revising one's inconsistent attitudes in any way, it matters not which? In any Non-exception, as *RN* observes, there is some revision that one has more reason to make than one has reason to remain with inconsistent intentions. For concreteness, imagine a Non-exception in which intending additionally to *Y* makes it less likely that one *X*'s, which one has Stronger Reason to achieve. It is clear what there is to be said for satisfying *C* by revising one's intention to *Y*; one has more reason to do that than one has to remain with inconsistent intentions. The difficulty is that one may *not* have more reason to satisfy *C* in other ways than one has to remain with inconsistent intentions. One may have more reason (to intend *X* and to intend *Y*) than (not to intend *X* and to intend *Y*) or more reason (to intend *X* and to intend *Y*) than (not to intend *X* and not to intend *Y*). And, whatever one intends, one may have conclusive reason to believe that if one *X*'s, then one does not *Y*. What, then, can be said equally in favor of these revisions?⁵¹ If nothing can be said equally in favor of these revisions, then *RN* would seem to capture the whole of the relevant truth.⁵² This problem is only intensified by Exceptions, in which there is not even one revision—one way of satisfying *C*—which one has more reason to make.

An instrumentally valuable disposition? One might suggest that, while satisfying *C* in any particular case may not lead one to attitudes that one has more reason to have in that case, a disposition to satisfy *C* over the long run does lead one to such attitudes over the long run.⁵³ First, even if there is something to be said for having a disposition to satisfy *C*, this does not, by itself, explain what there is to be said for satisfying *C*.⁵⁴ This is a problem well known to indirect consequentialists. From the fact that one is morally required to be disposed to conform to commonsense morality (because this disposition will lead one to do the most good over the long run), it does not follow that one is morally required to conform to commonsense morality in a particular case in which this would not do the most good.

Second, it is at any rate not obvious that a disposition to satisfy *C* leads one to attitudes that one has more reason to have over the long run. The basic dilemma is this: On its own, the disposition doesn't help, because it can just as soon send one into a worse state as a better one. And once we add enough other dispositions to make one reliably sensitive to reason, then the disposition becomes either superfluous, or, in Strong Exceptions, a liability.

To explain: The disposition in question cannot be a disposition, when one has inconsistent intentions, to adopt some specific consistent set of attitudes. Showing that a disposition to adopt some specific consistent set is valuable would not explain the normativity of *C*. For one does not manifest this disposition when one

adopts some consistent set *other* than the specified set, even though one still satisfies C. The relevant disposition—call it the ‘C-disposition’—must be a disposition, when one has inconsistent intentions, not to adopt some specific consistent set, but instead to adopt, in some sense, any consistent set ‘indifferently’. Only that would be a disposition that strove for consistency *as such*. It is hard to see what sense we can give ‘indifferently’, other than to say that the C-disposition is a disposition, when one has inconsistent intentions, to adopt the consistent alternative that one is otherwise least disposed to avoid, and if there is more than one consistent alternative that one is otherwise least disposed to avoid, to select from among these with equal probability.

Thus, the claim to be proved is that the C-disposition leads one, over the long run, to attitudes that one has more reason to have. The case for the claim starts in arrears. In Strong Exceptions, the C-disposition only leads to attitudes that one has less reason to have, and in other Exceptions, to attitudes that one has no more reason to have. In Non-Exceptions, there will be at least one way of satisfying C that leads one to attitudes that one has more reason to have. But there will often also be ways of satisfying C that lead one to attitudes that one has no more reason, or less reason, to have. Since the C-disposition selects ways of satisfying C with equal probability, there will be gains and losses in the Non-Exceptions. Why expect that the net effect of the C-disposition, taking Exceptions and Non-Exceptions together, will be positive?

‘When the C-disposition works together with some complementary disposition’, it might be suggested, ‘the upshot, in the Non-Exceptions, is not to select among *all* ways of satisfying C with equal probability, but instead to select among *only* those ways of satisfying C that one has *most reason to have*.⁵⁵ This gain in the Non-Exceptions from the addition of the C-disposition to the complementary disposition outweighs the losses in the Exceptions’.

But consider what this complement would have to be. It could not be simply the disposition to have the attitudes that one has most reason to have, since then the C-disposition would make no difference in the Non-Exceptions (and only a negative difference in the Exceptions). Nor could it be a disposition to have *consistent* attitudes that one does *not* have most reason to have. The C-disposition would again make no difference, and, at any rate, the two dispositions would not lead one to the attitudes that one has most reason to have in the Non-Exceptions. Instead, the complement would have to be a disposition (either to have the attitudes that one has most reason to have, or to have inconsistent attitudes). The C-disposition would then make a difference, by ruling out the second disjunct and steering one into the first.

The suggestion, when spelled out, is bizarre. In order to have the complement, one must already have the ability to *detect*—that is, to *respond differentially* to, whether consciously or unconsciously—which pattern (or patterns) of attitudes one has most reason to have. But if one has the ability to detect this pattern, then why not be disposed to respond to this detection *simply by adopting this pattern*? Why be disposed, instead, to respond to this detection *by either having this pattern, or having the inconsistent pattern*? Of course, the C-disposition would then rule out

the second disjunct and steer one into the first. But why bother with this detour in the first place? The *C*-disposition does nothing, it seems, unless other dispositions are rigged to give it busy work.

A rule for deliberation? It might be suggested that *C* helps one to deliberate. But, in fact, *C* does not itself tell one to deliberate in any particular way. It simply tells one to make one's intentions consistent as such. Nor, if one consciously sets out to satisfy *C*, is there much to deliberate about. After all, any revision satisfies *C* as well as any other.

This is not to deny that thoughts about consistency may be helpful in deliberation. First, if one's intentions are inconsistent, this may be *evidence*, in light of *RN*, that at least one of one's attitudes violates a requirement of reason. Thus, as I discuss in section 7, inconsistency may trigger a 'second-order' requirement of reason to deliberate about which attitude, if any, is against reason. Second, it might simplify one's deliberation to have a *policy*, when one believes that if one *X*'s, then one does not *Y*, not to consider one's reason (to intend to *X* and to intend to *Y*) and instead to assume that the upshot of such consideration would be that one does not have sufficient reason to intend both.⁵⁶ (Whether this simplification is a net benefit is a further question. If one must already consider the reason (to intend *X* and not to intend *Y*), (not to intend *X* and to intend *Y*), and (not to intend *X* and not to intend *Y*), in order to determine which of these options to take, it is not clear how much further cost is involved in considering the reason (to intend *X* and to intend *Y*). Most of the relevant factors are already in view. And there is a cost in not considering this reason: namely, that in the Strong Exceptions, one will overlook what is in fact a better option.⁵⁷) Yet, to repeat, *C* is neither a second-order requirement of reason, nor such a policy. *C* neither tells one to respond to inconsistency by trying to identify which attitude, if any, is against reason, nor tells one to avoid considering one's reason (to intend to *X* and to intend to *Y*), when one believes that the two are jointly unrealizable.⁵⁸

Facilitating autonomy? Recently, in important and ramifying work, Bratman has suggested that a special kind of intention—namely, 'a self-governing policy' about whether to treat considerations as reason-giving—may be constitutive of a central and valuable form of autonomy (as well as related goods, such as certain structures of social engagement).⁵⁹ Such policies constitute this form of autonomy only if they are uncontested, Bratman argues, since only then is one 'satisfied', in the sense of Frankfurt 1991, with them. And they are uncontested only if they satisfy *C*.⁶⁰ This suggests that the value of autonomy itself might provide reason, perhaps even conclusive reason, to satisfy *C*, at least where self-governing policies are concerned.

While attractive as an account of autonomy, this falls short of an explanation of the normativity of *C*, in two ways. First, it explains the normativity of *C* only insofar as *C* applies to self-governing policies, or to intentions that affect self-governing policies. But *C* applies to other intentions. Second, making self-governing policies consistent *as such* does not facilitate autonomy. If one achieves consistency by dropping *both* self-governing policies, one is no closer to autonomy. One is left without any self-governing policy at all, rather like the

‘wanton’ of Frankfurt 1971. This is not to deny, of course, that autonomy may be a reason to adopt *certain* consistent patterns of self-governing policies, just as Competition, Stronger Reason, and the like are often reason to adopt *certain* consistent patterns of ordinary intentions.⁶¹

A constitutive truth? Next, it might be suggested that some ‘constitutive’ or ‘conceptual’ claim underwrites the normativity of *C*. The claim that:

For an attitude to be an intention just is (in part) for it to satisfy *C*,

would not support, and may even be incompatible with, the normative claim that intentions *ought* to satisfy to *C*. The claim that:

For a subject to be an agent with intentions just is (in part) for certain of his attitudes to satisfy *C for the most part*.⁶²

would explain at most why we have reason to satisfy *C for the most part*, not why we have reason to satisfy *C in any particular case*. Lastly, consider:

For an attitude to be an intention just is (in part) for *C* to *apply* to it—so that to *know what* an intention is just is, in part, to *know that C* is a norm for intention,

and:

For a subject to be a (self-conscious, rational) agent with intentions just is (in part) for him to accept that *C* applies to certain of his attitudes—so that not to accept that *C* is a norm for intention is to cease to be a (self-conscious, rational) agent with intentions.⁶³

These claims do not answer the question being raised from the deliberative point of view, ‘Why should I *satisfy C*?’ At most, they answer the different question, ‘Why should I *accept that I should satisfy C*?’, forcing the deliberator to accept, awkwardly, *that* he should satisfy *C*, without any answer as to *why* he should. More generally, once we have distinguished *C* and *RN*, these constitutive claims seem far less plausible. Perhaps the real constitutive claims in the vicinity involve *RN*, instead of *C*. The conclusion of this paper, in effect, follows up on this suspicion.

An end in itself? Often when we cannot explain *in other terms* why we ought to satisfy a putative norm, the counsel of good sense is to acknowledge that we ought to satisfy the norm *for its own sake*. But it is not, as far as I can see, the counsel of good sense in this case. Simply put, it seems outlandish that the kind of psychic tidiness that *C*, or any other requirement of formal coherence, enjoins should be set alongside such final ends as pleasure, friendship, and knowledge.⁶⁴

Non-evidential reason for belief? A final problem, with all of the proposed reasons for satisfying *C* that we have considered, is that they are at odds with Evidentialism.⁶⁵ Reasons for satisfying *C* would be, to some extent, reasons for

revising the belief that if one *X*'s, then one does not *Y*. And yet none of the proposed reasons that we have considered is purely evidential.⁶⁶

A matter of evaluation, rather than deliberation? Notice we have been asking: What, if anything, can be said for satisfying *C*, within, or addressed to, the first-person standpoint of deliberation—the standpoint from which we decide what to believe or choose? We could ask a different question: How, if at all, can a person's satisfying *C* be positively appraised from the (typically) third-person standpoint of evaluation—the standpoint from which we approve or disapprove, praise or blame what someone believes or chooses?⁶⁷ Our failure to find an answer to the first question, it might be said, need not rule out an answer to the second. 'We can evaluate a person as beautiful, or some organ of hers as functioning properly, without implying that there are reasons, that might carry weight within the deliberative standpoint, for her to believe or choose anything. Likewise, we can evaluate someone as functioning properly, or manifesting a virtue, when she makes her intentions consistent, without implying that there are reasons for her to revise her attitudes as she does. This shift of focus might overcome some of the problems that we have encountered. While the fact that a disposition is good to have does not explain why one has *reason* to manifest it on any particular occasion, it might explain why manifesting it on any particular occasion merits some positive *appraisal*, as the expression of a virtue, or an episode of proper functioning. And while it may be the case that any *reason* for belief must be evidential, it does not follow that beliefs cannot be positively *appraised* on nonevidential grounds (e.g. as charitable, loyal, or consoling)'.

But still, and for much the same reasons, it is obscure why satisfying *C* should merit positive appraisal. The answer cannot be, as we have seen, that the *C*-disposition is worth having because it leads one to attitudes that one has more reason to have, or because it partly constitutes intention.⁶⁸ And, at any rate, an answer to this evaluative question would not be an answer to our original, deliberative one.⁶⁹

7. Completing the Error Theory for C: Satisfaction and Rationality

So far, according to the Noncognitivist Version of the Error Theory, *RN* explains *V*'. But it might be argued that there is something more to explain:

Satisfaction Claim (S): In many cases, one *satisfies* a requirement, that one would not have satisfied by persisting with inconsistent intentions, by making one's intentions consistent, *even in a way that does not lead one to attitudes that one has more reason to have.*

Suppose, for example, that one intends to *X* and intends to *Y*, for which one has conclusive reason. However, one believes, *without* sufficient reason, that one cannot both *X* and *Y*. Now contrast two responses to this belief. The first is simply to continue along with these inconsistent intentions, making no change. The second is to revise one of the intentions, say the intention to *Y*. 'By responding in

the second way', one might argue, 'doesn't one satisfy *some* requirement, or at least do well in *some* respect? And wouldn't one have failed to satisfy this requirement, or failed to have done well in this respect, if one had responded in the first way? *RN* does not explain this, since the second response does not take one to attitudes that one has more reason to have. Indeed, one only loses an attitude that one has reason to have. By contrast, *C* would explain this'.

To begin with, *S* does not say, as *C* would imply, that in *every* case, one satisfies some requirement, or does well in some respect, by revising one's intention to *Y*. For it is not clear that this is true. If one believes that one is in a Strong Exception, for example, then it is not clear that one satisfies some requirement, or does well in some respect, by revising one's intention to *Y*. Revising one's intention to *Y* might be simply akratic. It seems, then, that something other than *C* explains *S*.

Here is a suggestion. Usually, one will be in a position to know *RN* (although presumably not in such regimented terms). And, usually, if one has inconsistent intentions, one will be in a position to know that one has. Therefore, usually, if one's intentions are inconsistent, one will be in a position to know that one probably lacks sufficient reason for one of the relevant attitudes. If the matter is sufficiently important, reason will require one to try to determine which attitudes, if any, one lacks sufficient reason for. Of course, in responding to this 'second-order' requirement of reason, one might not arrive at a judgment that one lacks sufficient reason for some attitude. For example, one might arrive at the judgment that one is in an Exception. Suppose, however, that one does arrive at the judgment that one lacks sufficient reason for some attitude. Then one is rationally required to revise that attitude, as is codified by:

Believed Reason (BR): If one believes at *t* that one lacks sufficient reason for *A*, then one is rationally required at *t* to revise, or to refrain from forming *A*, going forward from *t*, on the basis of the content of this belief.⁷⁰

The suggestion, then, is that the requirement in *S*—the requirement that one satisfies by making one's intentions consistent, even in a way that brings one no closer to reason—is either the second-order requirement of reason, or the rational requirement *BR*. If one satisfies both of these requirements, then one will make one's intentions consistent. Nevertheless, one might satisfy both of these requirements without acquiring attitudes that one has more reason to have. This is because one might reach a false conclusion about which attitude reason requires one to revise. If this is what happens in our example, then one falsely concludes that one has more reason (to intend to *X* and not to intend to *Y*) than (to intend to *X* and to intend to *Y*), and so, in accordance with *BR*, revises one's intention to *Y*.

If this amendment to the Error Theory is to be viable, *BR* must avoid the problems that afflict *C*. *BR* avoids the Problem of Extension, because in cases in which one knows that one is in an Exception, one does not violate *BR* by keeping inconsistent intentions. *BR* also avoids the Problem of Conflict. It is true that *BR* will rationally require one to violate reason, if one has a *false* belief about what

reason requires. But this does not violate Ideal Compatibility, since the conflict could have been avoided, by not forming that belief.

It is less obvious, however, how *BR* avoids the Problem of Normativity. After all, the same considerations that show that we have no reason to satisfy *C* even in a way that leads us to attitudes that we have no more reason to have, show likewise that we have no reason to satisfy *BR* even in a way that lead us to attitudes that we have no more reason to have: that is, to satisfy *BR* even when our belief about our reasons is false. While this is all true, the original worry was that there was nothing that could be said, within or to the first-person standpoint of deliberation, for satisfying *C* even in the wrong way. And there is, in a way, something that can be said, within or to the first-person standpoint of deliberation, for satisfying *BR* in the wrong way. Take someone who is required by *BR* to revise the intention to *Y*. Given that *BR* requires her to revise this intention, she must satisfy its antecedent, by judging that she lacks sufficient reason for it, that reason requires her to revise it. Since she judges that reason requires her to revise it, and since revising it is what *BR* requires, it will seem to her, of what *BR* requires, that she is required to do that. It is not as though she sees herself as having, or needing, some *special* reason to realize such-and-such a pattern among her attitudes, like the reason for satisfying *C* for which we searched in vain. Instead, she just judges, perhaps falsely, that reason requires her not to intend additionally *Y*: say, that given that *X*-ing matters more, it would be foolish to risk failing to *X* by trying to *Y*. Furthermore, we can advise, or do something that looks like advise, her to believe as *BR* requires, by drawing her attention to the content of her own belief that reason requires her not to intend *Y*. In sum, while reason does not in fact require one to satisfy *BR*, we can explain why it will inevitably *seem* to one as though, of what would satisfy *BR*, reason requires it. And this may be enough for the error theory for *C* that we are trying to construct.⁷¹

A different worry is that this does not fully explain *S*. For it might be said that one sometimes satisfies a requirement by making one's intentions consistent in a way that leads one no closer to reason and satisfies neither the second-order requirement, nor *BR*. First, one can make one's intentions consistent *akratically*: in defiance of one's own judgment about the reasons for those intentions.⁷² Second, one might make one's intentions consistent *unreflectively*: without any conscious reflection on one's reasons at all.

Akratic and unreflective cases may involve the manifestation of an unconscious disposition. Where unconscious dispositions are at issue, there is no standpoint of deliberation, and so our normative question, 'Why conform in this way?' does not arise. However, manifesting an unconscious disposition may qualify one, or some subpersonal part of one, for positive appraisal. Consider an unconscious disposition not to have *A*, when in conditions in which reliably, but fallibly, one lacks sufficient reason to have *A*. First, because this disposition is reliable, it leads one closer to reason over the long run. This explains why its manifestations are appraised positively. Second, its manifestations make one's intentions consistent. Finally, because this disposition is fallible, its manifestations may lead one no closer to reason in particular cases. Such a disposition may

be manifested in an akratic or unreflective case. This would explain why even in some akratic and unreflective cases, the subject sometimes does well, in at least one respect.

There is, of course, no guarantee that such a disposition is manifested in every akratic or unreflective case. But it is not clear that this is a problem. For, again, it is not clear that we do well in *every* akratic or unreflective case. (Recall the akratic case in which one believes, rightly, that one is in a Strong Exception.) Perhaps we do, if we manifest an unconscious disposition of the kind just described. But otherwise it simply is not clear.

8. Conclusion

I have tried to extend Raz's approach to instrumental rationality to practical consistency.⁷³ If this extension succeeds, it suggests that the myth of instrumental rationality is a special case. The broader myth is of a set of principles that enjoin formal coherence as such. Instead, there is only the structure of reason, and the shadow that this structure casts, insofar as we are rational, on our beliefs and choices.

This broader myth has a more descriptive, or constitutive, version. On this telling, it is essential to beliefs and intentions that they strive for formal coherence as such. What it is for something to be a belief is, in part, for it to be disposed either to attract the beliefs that follow from it, or to retreat; and for it to be disposed either to repel beliefs that are inconsistent, or to give way to them. Likewise, what it is for something to be an intention is, in part, for it to be disposed either to muster intentions for the apparent necessary means to it, or, failing that, to withdraw from the field; and for it to be disposed either to repulse, or to make way for, intentions for what seems jointly unrealizable (i.e. the C-disposition). The claim that the constitutive 'aim' or 'function' of intention is to coordinate action, by way of these dispositions, might be seen as a version of this myth.

If we follow Raz's lead, I think, we find our way to a different view. Whatever we say about the essence of belief and intention, there are at least reasons for beliefs, which have to do with truth, and reasons for intentions, which have to do with what they are apt to bring about. Moreover, believers and intenders are, to a greater or lesser extent, sensitive to these reasons, either via their beliefs about them, or via unconscious mechanisms.⁷⁴ While it may be tempting to think that a tendency for formal coherence as such is somehow part of this sensitivity, this is an illusion. We saw this with respect to the C-disposition, and similar things can be said about analogous dispositions.⁷⁵ What it is true to say is that, as comments on reason patterns, such as *RN*, show, insofar as believers and intenders are sensitive to their reasons, their attitudes will, by and large, be formally coherent. This is what nourishes the myth. But this formal coherence among their attitudes will be just by-product, or epiphenomenon, of their sensitivity to reason. It will not be the result of any tendency to formal coherence as such.

To put the point another way, suppose we set out to 'construct' a creature with intentions. We give our creature the capacity for states (i) that, when they have the content, to X at t , lead it to X when it believes that t is now, (ii) that tend to be undermined by its belief that it will not X , (iii) that (like many other attitudes) are usually epistemically available to it, (iv) that (like many other attitudes) have an inertial tendency, other things equal, to persist once formed,⁷⁶ and (v) that (like many other attitudes) are sensitive to its beliefs about, or other unconscious monitorings of, its reason for such states (but that it can form even when, as in cases of underdetermination, it does not judge that it has conclusive reason for them). Finally, (vi) we give the creature a general sensitivity, whether conscious or otherwise, to its reason for such states. This must involve, at a minimum, sensitivities to the Facilitative Principle and to the way that its reason depends on what is likely to happen.

We must give our creature this much, if we are to make it recognizable as an intender. And if we give it this much, then, without building in anything further about coordination, anything about the C -disposition or kindred tendencies, we have succeeded, I think, in making it an intender. We have made our creature, give or take, us.⁷⁷ Why will its states usually be consistent? Because its general sensitivity to how changes in the likely future (e.g. changes in the weather, in its fellows' readiness to cooperate) change its reason will apply equally to changes in the likely future brought about, specifically, by its own such states. Why will its states usually be means-end coherent? Because of the same sensitivity and, more fundamentally, because of its responsiveness to the Facilitative Principle. No doubt, there is a connection between intention and coordination. But, if this line, inspired by Raz, is correct, this connection is more of the nature of a theorem than of a definition.⁷⁸

Niko Kolodny
Department of Philosophy
314 Moses Hall #2390
University of California, Berkeley
Berkeley, CA 94720-2390
USA
kolodny@berkeley.edu

NOTES

¹ See, for example, Broome 1999, 2005a; Wallace 2001; and Scanlon 2007.

² There are other putative rational requirements. I discuss Believed Reason below: the requirement of coherence between beliefs about our reason for attitudes and those attitudes themselves. There are also internalist requirements of epistemic justification: narrow-scope requirements of informal coherence between cognitive states (beliefs and perceptual states). A third class might be internalist requirements of practical justification: narrow-scope requirements of informal coherence among cognitive states with non-normative contents and intentions or desires. See Parfit, ms.; Setiya 2007a; Smith 2004b;

and Wedgwood 2003. While these last two kinds are 'rational requirements', in the sense of requirements governing relations among attitudes independently of anything beyond those attitudes, they stretch our ordinary attributions of 'irrationality', which are restricted to cases in which the subject is more immediately at odds with himself. It was this more common sense of 'rationality' that I mostly had in mind in Kolodny 2005, although I was unclear about this.

³ Compare too the rationalist position described by Smith 2004a: 250: 'If morality requires some limited form of altruism then ... the principle of limited altruism is a principle ... on all fours with *modus ponens* and *modus tollens* and the principle of means-ends'.

⁴ Raz 2005a. See also Raz 2005b.

⁵ This departs from the letter, but not the spirit of Raz's formulation: 'When we have an undefeated reason to take an action, we have reason to perform any one (but only one) of the possible (for us) alternative plans that facilitate its performance' (Raz 2005a: 6). I discuss some worries about this treatment of *ME* in Kolodny 2007.

⁶ See, for example, his discussion of requirements of logical consistency and closure on belief, in Raz 2005 a and b, as well as Raz 1999, ch. 4 n. 6. Scanlon 1997 and 2007; and Schroeder forthcoming, I believe, are drawn to similar views.

⁷ I read this belief as (i) if one succeeds in *X*-ing, then one does *not succeed* in *Y*-ing, rather than as (ii) if one succeeds in *X*-ing, then one *fails* in *Y*-ing. There is a substantive difference if there are conditional intentions that one can *not succeed* in without *failing* in. See Ferrero forthcoming. For example, one might not succeed in one's intention (to accept the job, if offered it), without failing in that intention, if one is not offered the job. In note 49 below, I discuss a case where the difference may matter.

⁸ See Bratman 1987; Broome 2007: 364; Harman 1976, 1986; O'Neill 1985; Searle 2001: 263. In a more general form, *C* might require one to avoid sets of any number of intentions that one believes one cannot jointly achieve. The phrase 'inconsistent intentions' is ambiguous as between the pair of intentions and the triple of the pair plus the belief that they are jointly unrealizable. Likewise, 'making one's intentions consistent' is ambiguous as between revising at least one of the pair and revising at least one of the triple. I hope what I mean will be clear from context.

⁹ See Bratman 1987, ch. 8.

¹⁰ See O'Neill 1985.

¹¹ See Gibbard 2003 and 2006 ms.

¹² McCann 1991. In turn, both McCann and I are heavily indebted to Bratman 1987, although we take a different moral from his insights.

¹³ This understanding is indebted to Hieronymi 2005.

¹⁴ Some might object: 'There cannot be reason to *intend* to *X*, even 'object-given' reason, which would derive from reason to *X*. This is because we cannot respond to our recognition of (would-be) reason for intending to *X*, even object-given reason, by intending to *X*'. But why think that we cannot so respond? First, one might argue that, in considering reason to intend, we would be deciding whether to decide. We would then be driven to deciding to decide to decide, and so on—an infinite regress of intentions. But this just equivocates on 'decide'. 'Deciding whether' to decide is not *intending* to intend, but instead *forming a judgment about one's reason* to intend. This does not invite a regress anymore than 'deciding' whether to believe something, where this is understood as considering one's reason to believe it. Second, one might argue that if we formed intentions in response to consideration of reason for intention, then forming intentions would itself be an action that would need to be controlled by a further intention. This

would threaten a genuine regress. But why should it follow from the fact that we form intentions in response to consideration of reason for intention that forming intentions is an action controlled by a further intention? It does not follow from the fact that we form beliefs in response to consideration of reason for belief that forming belief is an action controlled by a further intention.

¹⁵ On some these possibilities, reasons and rationality may overlap. In Kolodny 2005, I more or less defined 'reasons' as facts independent of attitudes. I am here defining 'reason' in a less committed way, as what matters in deliberation.

¹⁶ It is natural to assume that the relevant evidence is the evidence available to the deliberator. Yet, if an advisor believes that she has evidence which is better than the deliberator's and which suggests that there is not reason for the deliberator to believe, or intend, such and such, then she will contradict the deliberator's belief that there is such reason. This suggests that the relevant evidence is *not* the evidence available to the deliberator. In Kolodny and MacFarlane ms., drawing on the relativist semantics of the kind proposed by MacFarlane forthcoming, we might conjecture that the truth of a proposition about a person's reason to act or believe is relative to the evidence available to the person assessing the proposition.

Although I do not assume this in the paper, I am attracted to the view that evidence consists in facts, which are available in some way: e.g. by being perceived, believed, known, or there to be known.

¹⁷ 'Causally' is a placeholder for whatever we need to avoid in the case of Newcomb-style phenomena, in which the fact that one intends to *X* is evidence that one will *X*, but not because the intention will cause one to *X*.

¹⁸ Notice the asymmetry: according to Aimism, it is no reason *to* intend *X* that it makes one more likely to take some means with a side *benefit* (although, of course, it might be reason to intend *those means*). But it is reason *not* to intend to *X* that it makes one more likely to take some means with a side *cost*.

¹⁹ I assume Aimism only to avoid unnecessary controversy. I doubt that it is true. The main case for Aimism seems to derive from variants of the 'toxin puzzle' of Kavka 1983. In Kavka's example, although one is aware of the 'state-given' benefits of intending to drink the toxin, one cannot form that intention. The argument is, first, that this shows that one cannot respond to the recognition of state-given benefits by intending and, second, that if one cannot so respond, then such benefits cannot be reasons for intending. Compare Hieronymi 2006 and Shah forthcoming. While I am sympathetic to the second step, I am not convinced of the first. In the toxin case, there is already a sufficient explanation why one cannot form the intention: namely, that one believes that one would not carry it out. See Kolodny 2005: 525 n. 18. So the puzzle does not establish that one cannot respond to the recognition of state-given benefits by intending when one believes that one would carry out what one intended. Indeed, we seem to respond to state-given benefits in entirely ordinary cases. First, I may have reason to stop deliberating at *t* whether to *X*. Sometimes, a way of bringing this about is simply to intend at *t* to stop deliberating. But, sometimes, another way, and, sometimes, indeed, a more effective way, is to intend at *t* to *X*. I may have this reason to intend at *t* to *X*, even if I would be no less likely to *X* if I intended to *X* at some later time. Second, if intending facilitates coordination, then this too may be a state-given reason for intending. Suppose I have conclusive reason either (to *X* and to *Y*) or (not to *X* and not to *Y*), it matters not which. The thing to avoid is either (*X*-ing and not *Y*-ing) or (not *X*-ing and *Y*-ing), both of which are equally bad. Suppose that as things stand, the chances are even that I will *X*. The fact that by intending to *X*, I increase the chances that I will *X*, is not, on its own, any reason to *X*, since it does not increase the chances that I

either (X and Y) or (not X and not Y). Suppose, however, that since I can rely on myself to maximize expected value, I can rely on myself to Y , if the chances that I will X become greater than even. Then intending to X does increase the chances that I will either (X and Y) or (not X and not Y), from even to greater than even. So I do have reason to intend to X . This is not only because it increases the chances that the X component of (X and Y) obtains, but also because, by increasing those chances, it increases the chances that the Y component obtains as well. For related discussion, see Harman 1976, and Pink 1991. If Aimism is false, this means only that there are more Exceptions, as I observe in a later note.

²⁰ Why, Scanlon 2008 asks, is it wrong for me to buy rat poison with the intention of using it to kill my wife, but not wrong to buy it with the intention of using it to kill rats? One answer, he suggests, is that when I intend to kill my wife, I am in a position to know that my buying rat poison makes the murder of my wife more likely. I am, so to speak, making myself an accomplice to murder, just as I would be if I supplied someone else with rat poison, when I was in a position to know that he would likely use it to kill his wife. In ensuing discussion, however, Scanlon retreats somewhat from this claim. For criticism, see Kolodny forthcoming.

²¹ This is to say, against Standard Deontic Logic, that conjunction elimination is not permitted within the scope of a deontic operator. Compare Jackson 1985, and Jackson and Pargetter 1986.

²² See Kolodny forthcoming for an application of predictive significance to *ME*, which amplifies Raz's treatment.

²³ Because the relation between intending to X and intending to Y depends not on what one believes, but instead on what is likely, it is not a relation of rationality, unless we suppose that what is likely is solely a function of one's beliefs and perceptual states.

²⁴ The suggestion that one's intention to X can affect, in the presence of other factors, one's reason to intend to Y , might seem to undermine the substantive view, which I find attractive, that one's reason is a function of facts independent of one's attitudes. As I explain in Kolodny forthcoming however, the essence of that view remains viable.

²⁵ See Davis 1984; Grice 1971; Harman 1976 and 1986; Setiya 2007b; and Velleman 1989, 2000, and 2007. Audi 1973 claims that if one intends to X , one believes that it is more likely than not that one will X . Broome forthcoming b claims that if one *believes* that one intends to X , then one believes that one will X . Anscombe 1957 is claimed by both advocates and critics of Confidence.

²⁶ I discuss these facts and their applications in Kolodny 2007.

²⁷ They do not guarantee this Comment, for reasons related to the paradox of the preface. But one expects that it will hold for trios of beliefs of the kind involved in the following Comment. There may be a more significant issue here for a Cognitivist Error Theory of a generalization of *C*, which would apply to larger numbers of intentions.

²⁸ For a similar approach to *C*, see Velleman 2007. There have been several 'cognitivist' approaches to the principle of instrumental rationality, *ME*. See Broome 2002 and forthcoming b; Harman 1976 and 1986; Setiya 2007b; Velleman 1989, 2000, and 2007; and Wallace 2001. For criticism of these forms of cognitivism, see Bratman forthcoming a and b; Brunero 2005; and Raz 2005a. As Wallace 2001 and Bratman forthcoming a explain, cognitivist approaches to *ME*, unlike cognitivist approaches to *C*, do not require Confidence. Wallace relies on the weaker principle that if one intends to X , then one believes that it is possible for one to X . But he could rely, I think, on the even weaker principle, endorsed by Bratman 1987 and Mele 1989a, that if one intends to X , then one does *not* believe that one will not X (or, alternatively, that it is impossible for one to X).

²⁹ See Bratman 1987: 37–41. See also Davidson 1978; Holton 2008; McCann 1986, 1989, and 1991; and Wallace 2001.

³⁰ See Harman 1986; Mele 1989a and b. As Holton 2008 observes, this seems a more natural thing to say about the rescue case than about the bookstore case.

³¹ Compare McCann 1989 and 1991.

³² See, for example, Broome forthcoming b; Velleman 2007.

³³ Compare, again, McCann 1991.

³⁴ It is unclear whether it would. Granted, to say simply, 'I intend to X', will usually lead a hearer to believe that you believe that you will X. So, if you do not believe this, then a hearer may have grounds to complain. But, as Davidson 1978 argued against (somewhat surprisingly) Grice 1971, this may be a pragmatic phenomenon. It is not contradictory to say, 'I certainly do *intend* to X, but I am not sure that I will X. I have been known to forget such things'.

³⁵ There may also be mixed situations, in which it is uncertain whether one will need to decide. For example, one might wait to see whether the weather report rules out vacationing in the mountains or by the sea, knowing that, if the report is fair skies for both, one will have to decide oneself. For analytical purposes, however, we can leave these mixed situations to one side.

³⁶ Bratman 1987: 137–138.

³⁷ That is, (to intend to X and to intend to Y) rather than (to intend to X and not to intend to Y).

³⁸ For example, intending to rub one's belly may make it less likely that one pats one's head because it distracts one. This may be so even if one *can*, with enough concentration, both rub one's belly and pat one's head.

³⁹ As Non-intention-preventing situations show.

⁴⁰ I am indebted to Tim Scanlon for offering me this solution.

⁴¹ Bratman 1987, ch. 8.

⁴² See Makinson 1965.

⁴³ See Kolodny 2007.

⁴⁴ Recall that we are now assuming that Confidence is false, so that one can have all of these attitudes without having inconsistent beliefs.

⁴⁵ If Aimism does not hold, then there will be even more Exceptions to RN. Arthur Ripstein (p.c.) suggests (albeit to a different purpose) the following example: In order to keep his wits sharp, a prisoner plays chess against himself, which he knows will only achieve its intended purpose if he both intends to win as white and intends to win as black. He has sufficient reason (to intend to win as white and to intend to win as black), but these are state-given reasons, which do not derive from any reason he has to win as white, or to win as black. This does not obviously threaten the Non-cognitivist Version, however, since these cases may simply be further exceptions to V. It is not clear, for example, that the prisoner violates any requirement of reason or rationality.

⁴⁶ Bratman suggested this response in his SPAWN comments. He left the possibility open in Bratman 1987: 32.

⁴⁷ Compare Holton 2008 and McCann 1991.

Holton suggests that, in Video Game, one may have 'partial' rather than 'all-out' intentions. 'An intention to F is partial iff it is designed to achieve a given end E and it is accompanied by one or more alternative intentions also designed to achieve E' (41). This might encourage one to try to save C, understood as a claim about all-out intentions, by claiming that it is rational to have *partial* intentions in the Known Exceptions, but not *all-out* intentions. One problem is that there are Known Exceptions in which it seems strained

to describe one's intentions as being for alternative means to a single *E*. (One might both intend to be a successful lawyer and intend to be a successful jazz musician, waiting to see which career will take off.) On the other hand, if we say that in all such cases one's intentions *are* for alternative means to a single *E*, then they *cannot, by definition*, be all-out. The possibility of violating *C* is ruled out by definitional fiat.

Holton explores, and rejects, the possibility that intentions are graded along the dimension of effort, or resources, devoted to executing them. But, as Holton recognizes, this corresponds to no plausible requirement of formal coherence (50). Here is how I understand the problem. To simplify, suppose that *X*-ing and *Y*-ing have the same value. Then it might be suggested that one is rationally required to see to it that one devote resources to executing intentions to *X* and to *Y* in proportion to the degree to which devoting those resources raises the probability of success. In other words:

$$\begin{aligned} & [b(X|i(X)) - b(X|0(X))] * i(X) \\ & = [b(Y|i(Y)) - b((Y)|0(Y))] * i(Y), \end{aligned}$$

where $i(X)$ is the degree of resources that I intend to devote to executing my intention to *X* and $b(X|i(X))$ is the degree to which I believe that I will *X*, conditional on my intending to devote degree $i(X)$ of resources to executing my intention to *X*. But this is implausible. Suppose I know that the postage on the application to North State is \$1, that the postage on the application to South State is \$2, that the schools are equally good, and that that either is equally likely to admit me. It is not irrational of me to intend to pay for the required postage for both, thereby intending to devote greater resources to executing my intention to be admitted to *Y*. It is more plausibly irrational to be *unprepared* to devote as much to increase the probability that one *X*'s as one is *prepared* to devote to increase by the same degree the probability that one *Y*'s. For example, if I am prepared to pay \$1.50 to apply to *Y*, but unprepared to pay \$1.50 to apply to *X*. But this constraint on 'preparedness' is a constraint not on *intentions*, but instead on *preferences*, or on *beliefs* about value. Indeed, it applies to us whether or not we intend anything.

⁴⁸ Bratman 1987, 117–118.

⁴⁹ Suppose we understand the content of these conditional intentions such that, if one cannot hit target 1, then one neither succeeds nor fails in one's intention (to hit target 1, if one can). Suppose we also accept only the alternative interpretation of the belief in *C* that was suggested in note 7 above: if one succeeds in *X*-ing, then one *fails* in *Y*-ing. Then these conditional intentions would not violate *C*, since succeeding in one intention does not imply failing in the other. But this threatens to prove too much. Ferrero forthcoming raises the possibility that all intentions are of the form: to *X*, if one can. If so, then on the alternative interpretation, agents cannot violate *C* no matter what intentions they have. It becomes an empty constraint.

⁵⁰ Actually, this is a simplification. Strictly speaking, as Selim Berker points out to me, the disjunction should continue: '*or* reason requires one (either (to intend to *X* and not to intend to *Y*) or (not to intend to *X* and to intend to *Y*)) *or* reason requires one (either (to intend to *X* and not to intend to *Y*) or (not to intend to *X* and not to intend to *Y*)) *or* reason requires one (either (not to intend to *X* and to intend to *Y*) or (not to intend to *X* and not to intend to *Y*)) *or* reason requires one (either (to intend to *X* and not to intend to *Y*) or (not to intend to *X* and to intend to *Y*))'.

⁵¹ It might be objected: 'At least in the Non-exceptions, *RN* entails *C*, so long as disjunctions may be introduced within the scope of "required", as Standard Deontic Logic allows'. If disjunction introduction is permissible, then *RN* also entails an endless list of 'idle' principles like *C*#: One is required (either not to intend *X*, or not to intend *Y*, or not to

believe if one X 's, then one does not Y , or to dance the Hokey-Pokey) or $C\#\#$: One is required (either not to intend X , or not to intend Y , or not to believe if one X 's, then one does not Y , or to intend X , intend Y , and believe that if one X 's, then one does not Y). Following Føllesdal and Hilpinen 1971 and Wedgwood 2006, one might defend idle principles by offering a pragmatic explanation of their seeming oddity. It is never conversationally appropriate to assert $C\#$ or $C\#\#$, because we are always in a position to assert the more informative RN . But if C is merely an idle principle—if its only justification is that it follows, via disjunction introduction, from RN —then it is never conversationally appropriate to assert C either, because we are always in a position to assert the more informative RN . This indicates, as is anyway evident, that those who do endorse C do not view it merely as an idle principle. This might be because they *confuse* it with RN . Or it might be because they accept the Satisfaction Claim of section 7 (which implies that C has a status that, say, $C\#\#$ lacks). At any rate, the position that I am questioning is that C is something more than an idle principle. The view that it is nothing more is largely sympathetic to my skeptical position.

Next: 'Granted, making any one of these other revisions does not *suffice*, by itself, to give one a set of attitudes that one has more reason to have. But it is still a *necessary means* to having a set of attitudes that one has more reason to have. One therefore has conclusive reason to make any one of these other revisions'. First, it is simply not true that one has conclusive reason to take necessary, but insufficient, means to having a set of attitudes that one has more reason to have. One may have conclusive reason (to take necessary, but insufficient, means to having a set of attitudes that one has more reason to have *together with* other means with which they will be jointly sufficient). But taking necessary, but insufficient, means, on its own—which is what is being envisioned in this case—is pointless. (Compare our earlier example of booking a room, but not attending the conference.) Second, there is no helpful sense in which by making one of these other revisions—say, *not* believing (that if one X 's, then one does not Y), *not* intending X , and intending Y —one takes *means* to having a set of attitudes that one has more reason to have—say, believing (that if one X 's, then one does not Y), intending X , and *not* intending Y . Of course, a necessary condition of

believing (that if one X 's, then one does not Y), intending X , and *not* intending Y is

either (believing (that if one X 's, then one does not Y), intending X , and *not* intending Y); or (*not* believing (that if one X 's, then one does not Y), *not* intending X , and intending Y).

So one might say that in *not* believing (that if one X 's, then one does not Y), *not* intending X , and intending Y one takes necessary means to (believing (that if one X 's, then one does not Y), intending X , and *not* intending Y). However, another necessary condition is

either (believing (that if one X 's, then one does not Y), intending X , and *not* intending Y); or (believing (that if one X 's, then one does not Y), intending X , and intending Y).

So, by the same logic, in *continuing* to have inconsistent intentions—or, indeed, in doing anything at all—one takes necessary means in the same way. This cannot account for the normativity of C . Thanks to Ralph Wedgwood for suggesting this line of objection.

⁵² There are two other reasons why the proponent of C seems committed to there being something to be said equally in favor of these revisions. First, it is implied by the claim that there is something to be said for satisfying C . Raz 2005a and b argues for this compellingly, in face of objections from Broome 2005b. Second, C would otherwise provide no explanation of the Satisfaction Claim discussed in section 7 and so would forfeit the support it derives from providing such an explanation.

⁵³ This is suggested by Bratman 1987: 35, 43, 46, 52; and forthcoming a.; and Broome 2005a and forthcoming a. For further criticism of this approach, see Kolodny 2008.

⁵⁴ This point is compellingly made by Raz 2005b.

⁵⁵ Or more reason to have. But I focus on the simpler case.

⁵⁶ For discussion of policies, see Bratman 1987. The policy must be understood in such a way that it is not satisfied by preventing oneself from having the belief.

⁵⁷ This is not to deny that deliberation must be simplified somehow. It would place excessive demands on our finite attention for us to keep track our reason entirely consciously. We must instead rely on unconscious mechanisms, refined and extended by experience and habit, to put certain items on, and keep other items off, the deliberative agenda, when appropriate. Moreover, we must often rely on unconscious mechanisms to reach appropriate conclusions directly, bypassing conscious deliberation entirely. Yet once we recognize the full generality of our need for such unconscious sensitivities, it becomes implausible that we have, or need, a sensitivity that is somehow specifically linked to what we intend. After all, we must rely on unconscious sensitivities to what is likely to happen *in general* (e.g. to the weather, to prices, to our spouses' patience, etc.) in order to keep certain things off the deliberative agenda and to put other things on. It would be surprising if we had a separable sensitivity to what is likely to happen, *specifically, as a result of what we intend*. Indeed, it is not clear how such a separation could even be effected. Since we interact with a world that is not itself under our control, what is likely to happen as a result of what we intend depends, in part, on what is likely to happen independently.

⁵⁸ Even if *C* somehow did tell one to avoid considering one's reason, it would still be inferior to such a policy. If one did consider one's reason (to intend to *X* and to intend to *Y*) and concluded that reason required one (to intend to *X* and to intend to *Y*), then *C*—unlike the policy—would forbid one from making use of the fruits of that consideration. Notice, also, that any benefit would be vitiated by Bratman's distinction between intending and endeavoring. Although *C* would save one from considering one's reason (to intend to *X* and to intend to *Y*), it would not save one from considering one's reason (to endeavor to *X* and to endeavor to *Y*), which is all but the same thing.

⁵⁹ See Bratman 2007.

⁶⁰ See Bratman 2000 and 2005.

⁶¹ If we ask whether the *C*-disposition might be valuable because it facilitates autonomy, then we get much the same dilemma that we encountered earlier. On the one hand, if one is not disposed to respond to reason, then the *C*-disposition never leaves one worse off and sometimes leaves one better off as far as autonomy is concerned. However, it may leave one worse off in other ways. And, if one is not disposed to respond to reason, then one arguably cannot be autonomous in the first place. On the other hand, if one is disposed to respond to reason, then one will respond to one's reason to be autonomous by resolving inconsistency in self-governing policies in such a way as to achieve autonomy. The *C*-disposition will again be superfluous.

⁶² Compare Davidson 1985.

⁶³ Compare Korsgaard 1996.

⁶⁴ Compare Wedgwood 2003.

⁶⁵ Since none of the putative reasons for satisfying *C* is a reason *for* intention, none is at odds with Aimism.

⁶⁶ Margaret Gilbert and Kieran Setiya (p.c.) suggest a further, 'bootstrapping' problem with the suggestion that reason requires one to satisfy *C*. Plausibly, if reason (or at least practical reason) requires one (either to *A* or to *B*) and one cannot (or weaker: will not) *A*,

then reason requires one to *B*. Now consider a case in which one can (or will) revise neither one's intention to *X*, nor one's belief that if one *X*'s, then one does not *Y*. Then the assumption that reason requires one to satisfy *C* implies that reason requires one not to intend to *Y*. Yet, one can easily imagine such cases in which reason intuitively permits, or even requires, one to intend to *Y*. Compare Greenspan 1975; Schroeder forthcoming; and Setiya 2007b. One might reply, along the lines described in section 5, either that the (otherwise requiring) reason to satisfy *C* is overridden by the requirement of reason to intend to *Y*, or that what reason permits, or requires, is not to intend to *Y*, but instead (either to intend to *Y*, or to endeavor to *Y*). However, these replies are vulnerable to the objections raised in section 5.

⁶⁷ A related question is whether *C* represents a regularity to which we can appeal in giving 'rational explanations', typically from the third-person point of view, of the subject's responses.

⁶⁸ I believe that we can explain why satisfying *other* rational requirements, such as Believed Reason, and the internalist requirements of epistemic and practical justification, merits positive appraisal. For one thing, satisfying them manifests dispositions that lead one to attitudes that one has more reason to have over the long run, at least when joined with certain other dispositions. There are also less instrumental grounds. About Believed Reason, see Kolodny 2008; about practical justification, see Setiya 2007a.

⁶⁹ 'If satisfying *C* merits positive appraisal', it might be said, 'then this is an answer to the deliberative question. The reason why you should satisfy *C* is that it would be an episode of proper functioning, or display a virtue'. However, the fact that *A*-ing would display the virtue of kindness, modesty, courage, etc. is not typically itself reason (let alone conclusive reason) to *A*: a consideration that carries weight in deliberation whether to *A*, even deliberation of the most reflective sort. First, this will typically get the relevant value, and locus of value, wrong. What matters is not one's own kindness, for example, but instead another's relief from suffering. Moreover, with many virtues, being prepared to take them, on reflection, as reasons is incompatible with possessing them. See Moran 1993. Finally, whether we have reason to *A* typically depends on whether the situation is a certain way (e.g. whether our ministrations would relieve his suffering), whereas whether we display the virtue depends only on our *believing* it to be that way. None of this counts against the 'Reasons' thesis of Setiya 2007a: 'The fact that *p* is reason for *A* to *phi* just in case *A* has a collection of psychological states, *C*, such that the disposition to be moved to *phi* by *C*-and-the-belief-that-*p* is a good disposition of practical thought, and *C* contains no false beliefs'. This does not entail that the fact that *A*-ing would display a good disposition of practical thought is *itself* a reason to *A*. Setiya argues persuasively that neither (as I once thought) the left-hand, 'deliberative' side of the biconditional is explanatorily prior, nor that (as Smith 1994 seems to suggest) the right-hand, 'evaluative' side is prior. Instead, I am now inclined to think, different questions privilege different sides. When we ask what to choose, for example, what matters are the reasons in favor of possible choices. The fact that someone with the relevant virtue, if she knew the relevant particulars, would *A* matters insofar as it indicates that, as she would discern, there are some reasons or other to *A*. When we ask whether we are to be praised or blamed for *A*-ing, by contrast, what matters is whether we displayed the relevant virtue. Whether there were reasons to *A* is neither here nor there.

⁷⁰ This is actually an abbreviation. I discuss the full principle in Kolodny 2007 and 2008.

⁷¹ This is the 'Transparency Account' of Kolodny 2005. I need to say more to defend it. For careful and forceful objections, see Bridges forthcoming and Hussain ms.

⁷² I thank Alex Sarch and Nishi Shah, in particular, for pressing me on this.

⁷³ I discuss other extensions in Kolodny 2007.

⁷⁴ Indeed, 'reason' here is unnecessarily restrictive (even though 'reason' is already meant very broadly, e.g., as agnostic between 'value-based' and 'desire-based' conceptions). We could just as well speak of sensitivity to what is likely to achieve a certain list of weighted objectives, such as nourishment, reproduction, or the destruction of all that is good and holy. This would cover children, lower animals, and perverse agents, who seek precisely what they have least reason to do.

⁷⁵ See Kolodny 2008.

⁷⁶ This feature of stability, however, may itself be a by-product, or epiphenomenon, of its sensitivity to reason.

⁷⁷ In his SPAWN comments, Bratman suggested that such a creature would lack a unified practical standpoint. It is not clear to me why this should be so. First, my creature has a unified practical standpoint in the sense that he has a unified set of beliefs about reason, which govern what he intends. Second, while my creature treats what will happen as a result of what he intends as affecting his reason for other intentions, my creature does not treat it as somehow an obstacle thrown up by a foreign intelligence, or by brute nature. He understands that what will happen as a result of what he intends is, *unlike* the cooperation of others or the weather, under his control and answerable to his reasons. Finally, as I indicated earlier, if Bratman is right that we have reason to be autonomous, then my creature will care about and strive for its own autonomy.

⁷⁸ This paper has benefited greatly from discussions at colloquia at the University of Toronto and the University of Oxford; a BAFFLE workshop at the University of California, Berkeley; the 2006 Workshop on Practical Reason at the University of Leeds; SPAWN 2007 on Practical Reason at Syracuse University; and a Law and Philosophy Workshop at the University of Southern California. I owe special debts to Michael Bratman, who gave prepared comments at SPAWN; and to Selim Berker, Andrei Marmor and his workshop, Jay Wallace, and an anonymous reviewer for the *European Journal of Philosophy*, for written comments.

REFERENCES

- Anscombe, G. E. M. (1957), *Intention*. Oxford: Basil Blackwell.
- Audi, R. (1973), 'Intending', *Journal of Philosophy*, 70: 387–403.
- Bratman, M. (1987), *Intention, Plans, and Practical Reason*. Cambridge, MA: Harvard University Press.
- (2000), 'Reflection, Planning and Temporally Extended Agency', in Bratman 2007: 21–46.
- (2005), 'Planning Agency, Autonomous Agency', in Bratman 2007: 195–221.
- (2007), *Structures of Agency*. Oxford: Oxford University Press.
- (forthcoming a), 'Intention, Belief, Practical, Theoretical', in S. Robertson, J. Skorupski and J. Timmerman (eds), *Spheres of Reason*.
- (forthcoming b), 'Intention, Belief, and Instrumental Rationality', in D. Sobel and S. Wall (eds), *Reasons for Action*. Cambridge: Cambridge University Press.
- Bridges, J. (forthcoming), 'Rationality, Normativity and Transparency', *Mind*.
- Broome, J. (1999), 'Normative Requirements', *Ratio*, 12: 398–419.
- (2002), 'Practical Reasoning', in J. Bermúdez and A. Millar (eds), *Reason and Nature: Essays in the Theory of Rationality*. Oxford: Oxford University Press.
- (2005a), 'Does Rationality Give Us Reasons?', *Philosophical Issues*, 15: 321–37.

- (2005b), 'Have We Reason to Do as Rationality Requires: A Comment on Raz', *Journal of Ethics and Social Policy*, 1.
- (2007), 'Wide or Narrow Scope?', *Mind*, 116: 359–370.
- (forthcoming a), 'Is Rationality Normative?', *Disputatio*.
- (forthcoming b), 'The Unity of Reasoning?', in S. Robertson, J. Skorupski and J. Timmerman (eds), *Spheres of Reason*.
- Brunero, J. (2005), 'Two Approaches to Instrumental Rationality and Belief Consistency', *Journal of Ethics and Social Philosophy*, 1: 1–20.
- Davidson, D. (1978), 'Intending', reprinted in *Essays on Actions and Events*. Oxford: Oxford University Press, 1980.
- (1985), 'Incoherence and Irrationality', reprinted in *Problems of Rationality*. Oxford: Oxford University Press, 2004.
- Davis, W. (1984), 'A Causal Theory of Intending', *American Philosophical Quarterly*, 21: 43–54.
- Ferrero, L. (forthcoming), 'Conditional Intentions', *Noûs*.
- Føllesdal, D. and Hilpinen, R. (1971), 'Deontic Logic: An Introduction', in R. Hilpinen (ed.), *Deontic Logic: Introductory and Systematic Readings*. Dordrecht: D. Reidel.
- Frankfurt, H. (1971), 'Freedom of the Will and the Concept of a Person', *Journal of Philosophy*, 68: 5–20.
- (1991), 'The Faintest Passion', in *Necessity, Volition, and Love*. Cambridge: Cambridge University Press, 1999.
- Gibbard, A. (2003), *Thinking How to Live*. Cambridge, MA: Harvard University Press.
- (2006) (ms), 'Insight, Consistency, and Plans for Living', Tanner Lecture on Human Values.
- Greenspan, P. (1975), 'Conditional Oughts and Hypothetical Imperatives', *Journal of Philosophy*, 72: 259–76.
- Grice, H. P. (1971), 'Intention and Uncertainty', British Academy Lecture.
- Harman, G. (1976), 'Practical Reasoning', reprinted in *Reasoning, Meaning, and Mind*. Oxford: Oxford University Press, 1999.
- (1986), *Change in View*. Cambridge, MA: MIT Press.
- Hieronymi, P. (2005), 'The Wrong Kind of Reason', *Journal of Philosophy*, 102: 437–57.
- (2006), 'Controlling Attitudes', *Pacific Philosophical Quarterly*, 87: 45–74.
- Holton, R. (2008), 'Partial Belief, Partial Intention', *Mind*, 117: 27–58.
- Hussain, N. (ms.), 'The Requirements of Rationality'.
- Jackson, F. (1985), 'On the Semantics and Logic of Obligation', *Mind*, 94: 177–95.
- Jackson, F. and Pargetter, R. (1986), 'Oughts, Options, and Actualism', *Philosophical Review*, 95: 233–55.
- Kavka, G. (1983), 'The Toxin Puzzle', *Analysis*, 43: 33–6.
- Kolodny, N. (2005), 'Why Be Rational?', *Mind*, 114: 509–63.
- (2007), 'How Does Coherence Matter?', *Proceedings of the Aristotelian Society*, 107: 229–63.
- (2008), 'Why Be Disposed to Be Coherent?', *Ethics*, 118: 437–463.
- (forthcoming), 'Aims as Reasons', in Samuel Freeman, Rahul Kumar, R. Jay Wallace (eds), a *Festschrift* for T.M. Scanlon.
- Kolodny, N. and MacFarlane, J. (ms.), 'Ought: Between Subjective and Objective'.
- Korsgaard, C. (1996), *The Sources of Normativity*. Cambridge: Cambridge University Press.
- McCann, H. (1986), 'Rationality and the Range of Intention', *Midwest Studies in Philosophy*, 10: 191–211.
- (1989), 'Intending and Planning: A Reply to Mele', *Philosophical Studies*, 55: 107–10.

- (1991), 'Settled Objectives and Rational Constraints', *American Philosophical Quarterly*, 28: 25–36.
- MacFarlane, J. (forthcoming), 'Epistemic Modals Are Assessment-Sensitive', in Brian Weatherson and Andy Egan (eds), *Epistemic Modality*. Oxford: Oxford University Press.
- Makinson, D. C. (1965), 'The Paradox of the Preface', *Analysis*, 25: 205–7.
- Mele, A. (1989a), 'Intention, Belief, and Intentional Action', *American Philosophical Quarterly*, 26: 19–30.
- (1989b), 'She Intends to Try', *Philosophical Studies*, 55: 101–6.
- Moran, R. (1993), 'Impersonality, Character, and Moral Expressivism', *Journal of Philosophy*, 90: 578–95.
- O'Neill, O. (1985), 'Consistency in Action', in *Constructions of Reason: Explorations of Kant's Practical Philosophy*. Cambridge: Cambridge University Press, 1989: 81–104.
- Parfit, D. (ms.), Climbing the Mountain.
- Pink, T. L. M. (1991), 'Purposive Intending', *Mind*, 100: 343–59.
- Raz, J. (1999), *Engaging Reason*. Oxford: Oxford University Press.
- (2005a), 'The Myth of Instrumental Rationality', *Journal of Ethics and Social Philosophy*, 1.
- (2005b), 'Instrumental Rationality: A Reprise', *Journal of Ethics and Social Philosophy*, 1.
- Scanlon, T. M. (1997), *What We Owe to Each Other*. Cambridge, MA: Harvard University Press.
- (2007), 'Structural Irrationality', in G. Brennan, R. Goodin, F. Jackson and M. Smith (eds), *Common Minds: Themes from the Philosophy of Philip Pettit*. Oxford: Oxford University Press, 84–103.
- (2008), *Moral Dimensions: Permissibility, Meaning, Blame*. Cambridge, MA: Belknap Press.
- Schroeder, M. (forthcoming), 'Means-End Coherence, Stringency, and Subjective Reasons', *Philosophical Studies*.
- Searle, J. (2001), *Rationality in Action*. Cambridge, MA: MIT Press.
- Setiya, K. (2007a), *Reasons without Rationalism*. Princeton, NJ: Princeton University Press.
- (2007b), 'Cognitivism about Instrumental Reason', *Ethics*, 117: 649–73.
- Shah, N. (forthcoming), 'How Action Governs Intention', *Philosophers' Imprint*.
- Smith, M. (1994), *The Moral Problem*. Oxford: Blackwell.
- (2004a), 'Objectivity and Moral Realism: On the Significance of the Phenomenology of Moral Experience', in *Ethics and the A Priori*. Cambridge: Cambridge University Press.
- (2004b), 'The Structure of Orthonomy', *Philosophy*, supp. vol. 55: 165–194.
- Velleman, J. D. (1989), *Practical Reflection*. Princeton, NJ: Princeton University Press.
- (2000), *The Possibility of Practical Reason*. Oxford: Oxford University Press.
- (2007), 'What Good is a Will?', in A. Leist (ed.), *Action in Context*. Berlin: de Gruyter.
- Wallace, R. J. (2001), 'Normativity, Commitment and Instrumental Reason', *The Philosophers' Imprint*, 1.
- Wedgwood, R. (2003), 'Choosing Rationally, and Choosing Correctly', in S. Stroud and C. Tappolet (eds), *Weakness of Will and Practical Irrationality*. Oxford: Oxford University Press, 201–29.
- (2006), 'The Meaning of "Ought"', *Oxford Studies in Metaethics*, 1: 12–60.