

Means-end coherence, stringency, and subjective reasons

Mark Schroeder

Published online: 1 February 2008
© Springer Science+Business Media B.V. 2008

Abstract *Intentions matter.* They have some kind of normative impact on our agency. Something goes wrong when an agent intends some end and fails to carry out the means she believes to be necessary for it, and something goes right when, intending the end, she adopts the means she thinks are required. This has even been claimed to be one of the only uncontroversial truths in ethical theory. But not only is there widespread disagreement about *why* this is so, there is widespread disagreement about in *what sense* it is so. In this paper I explore an underdeveloped answer to the question of in *what sense* it is so, and argue that resolving an apparent difficulty with this view leads to an attractive picture about *why* it is so.

Keywords Intention · Instrumental rationality · Means-end coherence · Wide-scope · Bratman · Broome · Ewing · Setiya · Three-envelope problem · Reasons

1 Means-end coherence

1.1 The problem

Zach intends to do some action, *A*. And he believes that to do *A*, he must do *B*. Zach bears an interesting and important normative relationship to *B*. It is an action that he believes to facilitate his intended end, and something is going wrong, if he intends *A*, believes *B* to be necessary for *A*, has reflected clear-headedly on this fact, and yet still fails to intend to do *B*. But it turns out to be hard to say exactly what this relationship is, or what, exactly, is going wrong with Zach in this situation.

M. Schroeder (✉)
Philosophy, USC School of Philosophy, University of Southern California, Mudd Hall
of Philosophy, 3709 Trousdale Pkwy, Los Angeles, CA 90089, USA
e-mail: maschroe@usc.edu

It is not, for example, that Zach ought to do *B*, as one might hold on a naïve view. For example, suppose that *A* is the action of hiring an assassin to murder his wife, and *B* is the action of paying the assassin. It is not, plausibly, true that Zach ought to pay the assassin. Nor is it plausibly true that Zach ought to intend to pay the assassin. But there is still something wrong with Zach if he intends to hire an assassin to murder his wife, believes that paying the assassin is necessary in order to do this, and has no intention whatsoever of paying the assassin, no matter how long and clearheadedly he thinks about the matter.

Philosophers have constructed a number of fantastic ideas about what is going on in this case, instead. According to *Wide-Scopers*, Zach ought to do some *disjunctive* action—to either-intend-to-pay-the-assassin-or-not-intend-to-hire-him-to-murder-his-wife.¹ It is because Zach does not satisfy this disjunction, that there is something wrong with him.² According to others, there is nothing *normative* going astray with Zach, in the narrow sense of ‘normative’ in which it contrasts with ‘evaluative’—rather, there is simply a failure of function of some kind. It is a malfunction of agents to not intend the means they believe to be necessary, but it doesn’t reflect on them, personally; rather it is more like having a leaky heart valve.³ And according to yet another view, it is true that Zach ought to intend *B*, but not in the ordinary, practical, sense of ‘ought’. Rather, it is true in the *epistemic* sense of ‘ought’. Kieran Setiya motivates this view as the last resort—on the grounds that none of the other ways of describing what is going on in Zach’s case even makes sense.⁴

I think, however, that all of these ideas, in addition to having their own problems, pass over a very simple one which possesses great promise. On the view I’ll suggest, Zach *subjectively* ought to do *B*, and *subjectively* ought to intend *B*. In the remainder of part 1 I’ll motivate this idea’s initial promise. Then in part 2 I’ll explain an important challenge that it faces—the challenge which Setiya has argued is fatal to the view, although I will state the challenge in a way that does not depend on Setiya’s substantive views about reasons. In part 3 I’ll explain how we can get around this challenge, provided that we are willing to make a certain kind of assumption about the nature of intention. The resulting view has the virtue, moreover, of having the right structure to also give us an elegant and complete explanation of *why* this relationship holds between Zach and his intentions. Those will be the main positive ideas of the paper. Then in part 4, I’ll start to refine these ideas, showing how to make the required assumptions about intention more plausible, and consider the advantages for an important phenomenon about deliberation over time. Finally, in part 5, I’ll show how the resulting view is able to accommodate and explain the relationship between intentions and necessary means as a limiting case of the relationship between intentions and merely facilitating

¹ Possibly adding ‘-or-not-believe-that-paying-the-assassin-is-necessary-to-hire-him’, as well. See Schroeder (2004) for discussion.

² See, for example, Hill (1973), Gensler (1985), Wallace (2001), and Broome (1999, 2001).

³ See, for example, Bratman (1987, forthcoming) and Raz (2005a).

⁴ Setiya (2007b). I should note that several of the authors cited in fact hold some combination of these views.

means. Because this is a deep problem for other accounts, I'll suggest that the ease with which my account deals with it is collateral evidence that the account is on the right track.

1.2 Ewing's problem and a working hypothesis

Our problem has the following structure: something goes wrong under certain circumstances. And those circumstances depend on what the agent believes. It is the actions Zach *believes* to be necessary means to doing *A*, to which he bears the interesting and important normative relationship—not the actions which *are* necessary means to doing *A*. So in looking to understand what is going on in Zach's case, it should be fruitful to examine other cases in which something goes wrong in circumstances that depend on what she believes. And fortunately, there are a number of other important cases in ethical theory with this very structure. Here is one: Raul believes that there is a vast, international conspiracy whose primary objective is to trick him into playing hopscotch. Raul needs professional help. His need for help depends on his beliefs. Saul, who has no such belief, does not have Raul's pressing need for professional help. Something is going wrong, if Raul does not get help, but not wrong, if Saul does not get help.

The thing that goes wrong in Raul's case, however, does not seem to be very much like the thing that goes wrong in Zach's case, if he does not intend *B*. For example, an impartial and benevolent bystander with all relevant information about Raul's case would advise him to seek professional help.⁵ But an impartial and benevolent bystander with all relevant information would not advise Zach to intend to pay the assassin. She would advise Zach to give up on his plan to hire an assassin to murder his wife. So in looking for cases to compare Zach's to, we should avoid cases like Raul's—something different seems to be going on, in them.

There are other important cases, however, that seem to differ from Raul's in the same way that Zach's case differs from Raul's. For example, there is an important relationship between an agent and the things she *believes* she ought to do. Suppose that Yves believes that he ought to do *A*, but does not do *A*. Something is going wrong in such a case. And as in Zach's case and Raul's case, it is something that depends on Yves's beliefs. But it does not seem to depend on his beliefs in the same way that Raul's case does. As we observed before, an impartial and benevolent bystander with all relevant information about Raul's case would advise him to seek help. But an impartial and benevolent bystander with all relevant information about Yves's case would not necessarily advise him to do *A*. For example, Yves's belief that he ought to do *A* might be *false*. It might even be that Yves ought *not* to do *A*. In that case, a benevolent and impartial observer would surely advise Yves *not* to do it.

So Yves's case is structurally like Zach's case in many ways. In each case, something can go wrong, which depends on their beliefs, but it is not the sort of thing that goes wrong in Raul's case, or else does not depend on their beliefs in the same way. Moreover, in neither case is it true that they ought to do the thing such

⁵ Compare Schroeder (2007b).

that, when they do not do it, something is going wrong. This is an important point. It sounds platitudinous to say that you ought to do what you believe you ought to do—that's just to say that you ought to let your conscience be your guide. But it can't be in general true that for each thing anyone believes she ought to do, she ought to do it. For if it were, then everyone would be infallible about what she ought to do. Which is false. People are not, in general, infallible about what they ought to do. But there is still some interesting relationship between them and the things that they believe that they ought to do.

Saying exactly what this relationship is, is known as *Ewing's Problem*.⁶ A number of philosophers have drawn attention to its parallel structure to the problem about intentions with which we began.⁷ According to a Working Hypothesis that it would be worthwhile to investigate, the same sort of thing is going on in each case. That is the Working Hypothesis that I am going to investigate in this paper. I will be defending the view that the Working Hypothesis is correct. In the next two sections, I will defend a version of Ewing's answer to his own problem, in contrast with the competing wide-scope answer. Then in the remainder of the paper I will show how this answer can be extended to the case of means-end coherence, in a way that yields not just an answer to *in what sense* something is going wrong with Zach, but an answer to *why*.

1.3 The wide-scope answer

Wide-Scopers like Broome (1999) have also been keen on the parallel between the problem about intentions and Ewing's Problem. They claim that their strategy works for both. But it is a miserable answer to Ewing's Problem. The Wide-Scope view is that what Yves ought to do, is to either-do-A-or-not-believe-that-he-ought-to-do-A. This is why, they, say, something is going wrong with him if he believes that he ought to do A and doesn't do A. He is falling astray of this particular thing that he ought to do—the one that says that he shouldn't be that way.

The advantage of the Wide-Scope answer to Ewing's Problem is that it is supposed to explain why it does not follow that Yves ought to do A. All that he ought to do is to either-do-A-or-not-believe-that-he-ought. And from this, they claim, it does not follow that he ought to do A. Another perfectly good way of doing this disjunctive thing is to change his belief. As a result, Yves's belief that he ought to do A is not, after all, infallible, and hence we have a solution to Ewing's Problem. I think there are at least three very bad problems for this Wide-Scope answer. The first is that it is symmetric in a way that the duty of conscience is not. And the second two are both ways in which Wide-Scopers *are*, after all, committed to concluding that in at least some kinds of case, Yves's belief really is infallible—the very problem wide-scoping is motivated in order to avoid.

⁶ Ewing (1953) and Piller (unpublished).

⁷ For example, Gensler (1985) and Broome (1999). I discussed both kinds of case in Schroeder (2004), although I now think that I was wrong to say that my remarks about 'desire' there apply equally well to the case of intentions.

The first problem for Wide-Scoping is that it is symmetric. It doesn't distinguish between acting in accordance with your moral beliefs and adopting moral beliefs in accordance with your actions, and as a result it fails to distinguish between following your conscience and the distinctive vice of rationalization.⁸ Rationalization is the vice of changing your beliefs about what you ought to do, because you are not going to do it, anyway. According to the Wide-Scope view, this is precisely as good a way of satisfying this requirement as is actually paying attention to what you believe and acting accordingly. If the thing that is going right about Yves when he believes that he ought to do *A* and as a result, does it, is also going on the case of the Ultra-Rationalizer, who never lets her actions be affected by her beliefs about what she ought to do, but rather immediately changes her mind about whether she ought to do them, then it seems to me that we have not quite captured what is going wrong in Yves's case.

A naïve response to this objection holds that though changing his beliefs is not ruled out by the instrumental principle, it *is* ruled out by some *other* principle governing *theoretical* reason, which says not to change your beliefs about what you ought to do, or some such thing. But this is short-sighted. In general, if Yves ought to do either *A* or *B*, and ought not to do *B*, then it follows that Yves ought to do *A*. This means that if there really is a principle according to which Yves ought not to change his belief, then in any circumstance in which that principle applies, the only way for him to fulfill his Wide-Scope requirement to either-do-*A*-or-not-believe-that-he-ought is for him to do *A*. And so it follows that in those cases, he ought to do *A*. But that just means that in any case in which Yves ought not to change his belief, it follows from the Wide-Scope view that his belief must be true. So the Wide-Scope view is still committed to holding that Yves's belief about what he ought to do is infallible in any case in which he is believing rationally. And this looks like a second bad problem for the Wide-Scope view.

Moreover, it follows from another plausible general principle that things are even worse.⁹ As Patricia Greenspan has pointed out, it follows from the fact that Yves ought to either do *A* or do *B* and that Yves *cannot* do *B*, that Yves ought to do *A*.¹⁰ But plausibly, in at least many cases, changing his beliefs about what he ought to do is not something that Yves can do as instantaneously as he can act on them. Even if he has the power to change those beliefs, doing so takes some time. Consider, then, the intervening time. That is time during which the only way that Yves can either do *A* or not believe that he ought, is to do *A*. So it follows from the Wide-Scope view and this principle that during that time, he ought to do *A*. But with respect to all of our beliefs that we are not already trying to change, we are *always* in this intervening time. Hence, it follows that people are in general infallible about what they ought to do, as long as they do not try to change their minds. Whenever they believe they ought to do something and are not already trying to change their mind

⁸ See Schroeder (2004) for the statement of this objection. Kolodny (2005, 2007) also presses a version of this kind of objection against a kind of Wide-Scope view.

⁹ Setiya (2007b) also advances a version of this argument, although mine requires weaker assumptions, because it takes into account time. Full credit for the argument goes to Greenspan (1975).

¹⁰ Greenspan (1975).

about it, what they believe is true—they really ought to do it. Again, to accept this conclusion is to have failed to solve Ewing’s Problem in the first place.

These, I think, are each very bad problems for Wide-Scoping as an answer to Ewing’s Problem. So if the hypothesis that the problem about intentions and Ewing’s Problem should receive similar answers is a good one, that means that Wide-Scoping can’t be right about intentions, either. (And indeed, the Wide-Scoping account of the case of intentions faces problems in its own right.¹¹) It is Ewing’s own answer to his Problem that gives us a more promising way of developing our Working Hypothesis.

1.4 Subjective oughts

Ewing’s own answer to his Problem is that ‘ought’ is ambiguous in English, having two important senses. It is true, Ewing held, that you ought to do what you believe you ought to do. But this does not mean you are infallible about what you ought to do, because the first ‘ought’ has a different meaning from the second. There is an *objective* sense of ‘ought’, and you have beliefs about what you ought to do in that sense. But there is also a subjective sense of ‘ought’, and it is true that you subjectively ought to do those things that you believe that you objectively ought to do.

Ewing also had a simple explanation of why this last thesis is true. He held that the subjective sense of ‘ought’ just *meant* ‘believes she objectively ought’, making the requirement to follow your conscience simply definitional of the subjective sense of ‘ought’. We needn’t go in for this idea of Ewing’s about how to understand the relationship between the two senses of ‘ought’, however, in order to see the attractions of his solution.

Moreover, it is a virtue of this solution that Ewing’s Problem is only one kind of case which warrants distinguishing a subjective sense of ‘ought’. In Ewing’s cases, someone has a belief directly about what she ought to do, but which is mistaken. But in many cases, agents have no explicit beliefs about what they ought to do. Still, even in those cases, there are things that they ought to do. And the things that they ought to do sometimes depend on features of their situation. In some situations we ought to do some things, while in other situations we ought to do others. And finally, the very features of our situations on which what we ought to do depend, can sometimes be features of which we are unaware, or even have false beliefs about.

Bernard Williams’ gin and tonic case is such a case.¹² We are to imagine that Bernie is in the sort of situation in which the thing for him to do, if his glass contains gin and tonic, is to take a sip. If his glass contains gasoline, on the other hand, he ought not to take a sip. So what he ought to do depends on whether his glass contains gin and tonic or gasoline. We are also to imagine that Bernie believes that

¹¹ I do think that the problems I’ve pursued are *more pressing* in the case of Ewing’s Problem, which is part of why I think it generates leverage to tackle the problem indirectly, by way of our Working Hypothesis.

¹² Williams (1981). See Schroeder (forthcoming) for extensive discussion.

his glass contains gin and tonic, but that in reality, it contains gasoline. Philosophers disagree about what Bernie ought to do. According to some, he ought to take a sip, for it is the only rational thing for him to do—the best option, given his beliefs. According to others, he ought not to take a sip, since gasoline is toxic. But very plausibly, according to a generalization of Ewing's solution, these philosophers are simply cottoning on to different ways in which we can use the word 'ought'. Sometimes we can use it in the objective sense, in which he ought not to take a sip. But it also has a subjective sense, in which he ought to take a sip. Both sides are right—about something. For there are two senses of 'ought'.

In truth, for our purposes it doesn't matter greatly whether there are really two senses of 'ought'. It doesn't really matter whether or not it is really okay to use the word 'ought' for both of these relations—though if it were, that would yield the best explanation of why 'you ought to follow your conscience' makes perfect sense. What matters to me is simply to point out that there is an interesting normative relationship of some kind between Bernie and taking a sip, and another interesting normative relationship between Bernie and not taking a sip. Both relationships are interesting, and we appeal to them for different purposes. For example, our practices of advice hinge much more heavily on the latter relationship, and our practices of blame and praise depend much more heavily on the former. I'm going to call them the subjective ought and objective ought relations, but that is just to give them names. If you think one or the other doesn't deserve to be called 'ought', I'm happy to defer, and you may call it something else.

The important point is that once we get past the obstacle of thinking that the interesting claim here is semantic, I think the claim that there is a subjective 'ought' leaves little about which to be alarmed. Everyone should agree that there is an interesting normative relationship between Bernie and taking a sip. My hypothesis about Ewing's Problem is that the same relationship holds between someone who believes that she ought to do something, and what she believes that she ought to do. This is the 'in what sense' answer to Ewing's Problem.

2 A first pass

2.1 The explanation of Ewing's problem

So far, I've isolated an answer to *in what sense* something is going wrong in Yves's case, but I haven't explained *why*. Ewing had a simple explanation of why. According to Ewing, 'Yves subjectively ought to do A' just *means* 'Yves believes that he objectively ought to do A'. That claim about what 'subjectively ought' means makes it trivial that you subjectively ought to do what you believe you objectively ought to do. But it doesn't accommodate Bernie's case. Bernie doesn't have any explicit beliefs about what he ought to do, but there is still something that he subjectively ought to do.

Reflection on Bernie's case, however, motivates a more general principle about the relationship between objective and subjective oughts, that provides an alternative to Ewing's explanation. What is distinctive of Bernie's case, after all,

is not that he believes that he objectively ought to do something, but that he believes something which, if true, entails that he objectively ought to do it. In general, the following seems¹³ like a plausible principle:

subjective ought test *X* subjectively ought to do *A* just in case *X* has some beliefs which have the following property: the truth of their contents is the kind of thing to make it the case that *X* objectively ought to do *A*.¹⁴

The subjective ought test is a general principle about the relationship between objective and subjective oughts. It is motivated by considering a range of cases that are more general than the pure Ewing's Problem cases like Yves's—including Bernie's. But it still provides us with an explanation of why Yves subjectively ought to do *A*.

This is because Yves's case passes the subjective ought test perfectly. In Yves's case, what he believes is that he objectively ought to do *A*. But if that is true, then it *is* the case that he objectively ought to do *A*. So if the subjective ought test is true, it is a general principle which can explain why Yves subjectively ought to do *A*:

P1 Yves believes that he objectively ought to do *A*.

P2 If this belief is true, then Yves objectively ought to do *A*.

P3 If Yves has some beliefs which, if true, entail that he objectively ought to do *A*, then Yves subjectively ought to do *A*.

C Yves subjectively ought to do *A*.

The explanation is simple. Premise 1 simply describes Yves's case: he believes he ought to do something. Premise 2 is trivial. And Premise 3 is the right-to-left direction of the subjective ought test. So if the right-to-left direction of the subjective ought test is true, then Yves subjectively ought to do *A*. I take this to be an explanation, because it subsumes Yves's case under a much more general principle—one which explains what is going on in Bernie's case, as well.

So where does that leave us? Recall that our Working Hypothesis was that what is going on in Zach's case is much like what is going on in Yves's case. And I've just shown how what is going on in Yves's case can be explained by appeal to general principles which apply to an even wider range of cases. If our Working Hypothesis is right, then we should conclude that Zach subjectively ought to do *A*, and if we want to understand *why*, we should look, as in Yves's case, to our general principles about how subjective oughts work. But this leads to a problem.

¹³ In Sect. 5 I'll consider and endorse an argument that this principle can't be exactly right, and show how to introduce refinements which get around this problem, but for now it is close enough. Certainly it has been widely accepted—for example, by Parfit (forthcoming), who calls the distinction between objective and subjective oughts the difference between reasons and rationality.

¹⁴ Sometimes tests like this one are formulated in counterfactual terms. It is easy to see, though, that the counterfactual test can't really be quite right. For example, it could be that the closest world in which Bernie's glass really contains gin and tonic is one in which he has promised to give up on drink, or satisfies some other condition which would make a difference to what he objectively ought to do.

2.2 The hitch

We can explain Yves's case by appeal to the subjective ought test, because the belief on which what Yves subjectively ought to do depends, is such that if it is true, then Yves objectively ought to do *A*. The hitch is that Zach's case is not like that. The belief on which Zach's situation depends is his means-end belief: that to do *A* he must do *B*. But that belief can be true, even though it is not the case that Zach objectively ought to do *B*. For it may simply not be the case that he objectively ought to do *A*. To return to our example, it is not too hard to imagine that Zach's belief is true, that to hire an assassin to murder his wife, he must pay the assassin. But it is still clearly not the case that Zach objectively ought to pay the assassin.

Moreover, this doesn't merely present an obstacle to our *explaining* why it is that Zach subjectively ought to do *B*. It presents an obstacle to its even being *true*. For according to the left-to-right direction of the subjective ought test, Zach subjectively ought to do *B* *only if* he has some beliefs which, if true, guarantee that he objectively ought to do *B*. But all that we know about Zach is that he intends to do *A* and believes that to do *A* he must do *B*. And it is easy to imagine cases in which this belief is true, but it is not the case that Zach objectively ought to do *B*.

In fact, we can raise the same problem by appeal to a principle that is much weaker than the subjective ought test. Even if the subjective ought test turns out to be false, the following principle is still highly compelling:

very weak ought test If *X* is completely opinionated about every factual question, has a complete credence (=1) in every proposition she believes, and all of her beliefs are true, then if *X* subjectively ought to do *A*, it follows that she objectively ought to do *A*.¹⁵

Now imagine that Zach is completely opinionated about every factual question, has a complete credence in every one of those beliefs, and that all of his factual beliefs are true. And imagine that he intends to hire an assassin to murder his wife and believes, truly, that paying the assassin is necessary to hire her. Then from the supposition that Zach subjectively ought to intend to pay the assassin, it would follow that Zach objectively ought to pay the assassin. But that seems false. Even in this case, there is something going wrong with Zach if he does not intend the means. But not even in this case is it plausible that Zach objectively ought to intend the means.

The problem is that our examples suggest that there is some close connection between subjective and objective oughts which is mediated by the beliefs that the subjective oughts depend on. And in the case of means-end coherence, Zach's situation seems to depend on his means-end belief. But there does not appear to be a connection between subjective and objective oughts in Zach's situation that is mediated by this belief. So it is hard to see how it could be true that Zach subjectively ought to intend the means.

¹⁵ The qualifications in the very weak ought test address a class of proposed counterexamples to the subjective ought test that I will discuss in part 5.

This, in essentials, is Kieran Setiya's argument against the idea that Zach subjectively ought to intend to pay the assassin, in his forthcoming paper, 'Cognitivism About Instrumental Reason'.¹⁶ Setiya claims that this view is no better than the view that Zach objectively ought to intend to pay the assassin, because it entails it, anyway. Which is just what we've seen follows, given that Zach's belief is true, from the subjective ought test. Setiya doesn't use the subjective ought test; he appeals to his own substantive analysis of reasons in terms of good practical reasoning.¹⁷ But I think that the problem I've outlined here is just a generalization on his argument in that paper, appealing to weaker assumptions.

2.3 How bad

The problem, I think, is not a trivial one for our hypothesis. It tells us that even the view about subjective 'oughts' is committed to the objectionable conclusion about what Zach objectively ought to do that we initially sought to avoid, at least in the special case in which Zach's means-end belief is true. I have elsewhere advocated a very general strategy for dealing with such unintuitive results,¹⁸ so I want to say briefly why I think that strategy would not work, here.

The strategy I have employed elsewhere involves weakening the claim that Zach objectively *ought* to intend to pay the assassin, to the weaker claim that there is an objective *reason* for Zach to intend to pay the assassin.¹⁹ The next step of the strategy is to distinguish between this bare existential claim, and the further thesis that it is a particularly weighty reason.²⁰ The strategy is to deny, in the cases of unintuitive reasons, that they are very weighty at all, insisting that they are of particularly low weight, and trying to explain why.²¹ And finally, the strategy is to explain why it seems like there is no reason at all for Zach to intend to pay the assassin, by explaining why pragmatic factors dictate that our negative existential intuitions about reasons are likely to be systematically misleading, in the case of reasons of very low weight, and to provide independent evidence that this is so.²²

Although I have never committed in print to any views specifically about the norms governing intentions, for some time I did wonder whether something like this strategy could be an important part of solving the problem of what is going on in Zach's case. There are a variety of reasons why I think this can't be right. But the principal one, is that given a generalization of our test for subjective 'oughts', it fails miserably to account for the *stringency* of what goes wrong with Zach when he fails to intend the means he believes to be necessary to his intended end.

¹⁶ Setiya (2007b).

¹⁷ See Setiya (2007a).

¹⁸ Schroeder (2004, 2005, 2007a, b).

¹⁹ Schroeder (2004, pp. 344–345).

²⁰ Schroeder (2005, 2007a), and especially (2007b, Chap. 5).

²¹ Schroeder (2007a) and especially (2007b, Chap. 7).

²² Schroeder (2005, pp. 6–11, 2007a, b, Chap. 5).

The distinction between objective and subjective reasons can be made in the same ways as that between objective and subjective ‘oughts’. It can be done by considering Ewing-style cases, or by focusing on Bernie’s case.²³ Because they are motivated by similar cases, similar tests seem to apply:

subjective reason test *X* has a subjective reason to do *A* just in case she has some beliefs which have the property, if they are true, of making it the case that *X* has an objective reason to do *A*.

Moreover, it seems plausible to suppose, in connection with the subjective reason test, that if the objective reason would be weightier, then the subjective reason is weightier. At any rate, I’ll work with this idea.

Here is the problem: the interesting normative relationship between Zach and intending to pay the assassin is a *stringent* one, a *requiring* one. It is very strong. Something goes *very* wrong with Zach, if he clearheadedly intends to hire an assassin, recognizes that hiring an assassin requires paying her, and has no intention whatsoever to pay her. So by our test, something should be going *very* wrong in the objective sort of way with Zach, if he intends to hire an assassin and hiring an assassin really does require paying her, and he has no intention of paying her. So even if we could accept that there is something just a little bit wrong with Zach in that case, because there is *some* objective reason for him to intend to pay the assassin, but only a very weak one—so weak that our intuitions about it might be misleading—that wouldn’t be enough. The test requires that there must be something objectively *stringent* going wrong with Zach in the case in which his belief is true, in order for there to be something subjectively stringent going wrong with him, when he has that belief.

The stringency of the connection between Zach and intending to pay the assassin magnifies, I think, the problem posed in the last section. By the subjective ought test and the assumptions that we are making so far, it seems not only to follow from the view that Zach’s case is like Yves’s and Bernie’s cases in that there is some objective reason for Zach to intend to pay the assassin, if his belief is true. It seems to follow that in order to account for the stringency of the requirement for means-end coherence, this view would be committed to holding that Zach is objectively *required* to intend to pay the assassin. And that is something that we can all agree is patently false. It is not the kind of thing that the strategy I mentioned is capable of explaining away. So we need, I think, another solution.

3 A positive view

3.1 Oughts and transmission

To see how to solve this problem, I think we need to step back and observe a very general fact about the transmission of oughts, which we can use to construct a generalization of Yves’s case. This generalization will allow us to see which

²³ See Schroeder (forthcoming).

assumption that we were implicitly making was leading to all of the fuss, and consequently which assumption we can make about intentions that will remove the problem. The fact about the transmission of oughts is simple:

ought transmission If X objectively ought to do A , and to do A X must do B , it follows that X objectively ought to do B .

In part 5 I'll generalize this principle to reasons, but to get the basic structure of the idea, that is all that we need for now. The idea behind this principle is that the force of the first 'ought' *transmits* to the second, by means of the necessary connection between them.²⁴ I'm going to assume that this principle is true for objective oughts.

If it is, then it follows from the subjective ought test that something similar is true for subjective oughts. Suppose that Xera subjectively ought to do A and *believes* that to do A she must do B . Since she subjectively ought to do A , by the left-to-right direction of the subjective ought test there must be some things she believes, such that their truth would guarantee that she objectively ought to do A . Call the set of those beliefs S . But Xera also believes that to do A she must do B . So consider this belief, along with those in S . If all of *those* beliefs are true, then the transmission principle, above, guarantees that Xera objectively ought to do B . So since Xera believes these things, it follows from the right-to-left direction of the subjective ought test that she subjectively ought to do B . So this induces a *reflection* of the transmission principle for objective 'oughts' for the case of subjective 'oughts':

transmission reflection If X subjectively ought to do A , and believes that to do A she must do B , it follows that X subjectively ought to do B .

So now consider Xera's case more carefully. Let me stipulate that the content of S is her belief that she objectively ought to do A . This makes her case a generalized version of Yves's case. Yves's case does not depend on any other beliefs that he has. He simply has the belief that he ought to do A , and as a result, subjectively ought to do A . But Xera's case is more general. She has a belief about what she ought to do, and she subjectively ought to do this, just like Yves. But due to the reflection of the transmission principle, she also subjectively ought to do other things, which she believes are necessary means to A .

But Xera also has an interesting property: even if her belief that B is necessary for A is true, that is not enough to guarantee that she objectively ought to do B . For it could be that her belief that she ought to do A is false. So Xera fails the same test as Zach did. She definitely subjectively ought to do B —that follows from transmission reflection. But given the transmission principle, this depends on more than one of her beliefs. In order to test it, we must look at *all* of her beliefs that it depends on, not just the transmitting beliefs. It depends both on her belief that B is necessary for A *and* on her belief that she ought to do A . And if *both* of those beliefs are true, then the transmission principle does guarantee that she objectively ought to do B .

²⁴ Compare Darwall (1983, p. 16), Raz (2005a, pp. 3–9).

3.2 A motivated conjecture

Xera's case leads to a motivated conjecture: that the same thing might be going on in Zach's case. We had good cause to hypothesize, after all, that Zach subjectively ought to intend to pay the assassin. His case had a lot in common with those of Yves and Bernie, because the thing that was going wrong with him when he didn't so intend seemed to depend on his beliefs in a very similar way—quite different from how the thing going wrong with Raul depended on his beliefs. Yet Zach's case still differed from those of Yves and Bernie—it differed from them in exactly the same way that Xera's case differs from them. So it is more than natural to wonder whether Zach's case isn't to be explained in the same sort of way as Xera's is.

If so, then we could explain exactly where Setiya's reasoning went wrong. It went wrong in supposing that Zach's means-end belief is the *only* belief that his subjective 'ought' depends on. This assumption leads to the wrong conclusion about Xera's case, so if Zach's case is like hers in the right way, then it would lead us astray in Zach's case in the same way. In Xera's case, of course, the relevant difference is that her subjective 'ought' depends on her means-end belief because it plays the role of *transmitting* a further subjective 'ought'. So if Zach's case is like that, then his mean-end belief must also simply be transmitting a further subjective 'ought'.

At a first pass, then, this leads to the idea of explaining Zach's case in this way:

P1 Zach intends to do *A*.

P2 If Zach intends to do *A*, then he subjectively ought to do *A*.

P3 Zach believes that to do *A*, he must do *B*.

P4 Transmission Reflection: if Zach subjectively ought to do *A* and believes that to do *A* he must do *B*, then he subjectively ought to do *B*.

C Zach subjectively ought to do *B*.

Premises 1 and 3 simply describe Zach's case. Premise 4 is the transmission principle for subjective oughts, which is both independently plausible and derivable (as I showed in Sect. 3.1) from the transmission principle for objective oughts and the subjective ought test. And the conclusion follows validly from the premises. So the only premise in need of further defense is premise 2.

So far, I've offered no explanation or defense of premise 2, so this doesn't yet amount to an explanation of why Zach subjectively ought to do *B*—only a model for one. But I hope to have argued that it is a *promising* model, and hence that premise 2 is worth trying to explain. The explanatory model that I am advocating answers the question of *in what sense* there is something wrong going on in Zach's case, if he intends to do *A*, believes that doing *B* is necessary for *A*, and fails to do *B*. This sense, is that he is not doing what he subjectively ought to do. Moreover, the explanation on offer answers Setiya's concern—it does not turn out, on this view, that Zach objectively ought to do *B*, even if his means-end belief is true. And finally, it potentially leads to an attractive explanation of *why* Zach subjectively ought to do *B*, which appeals, in addition to whatever assumptions we need in order to explain

premise 2, only to plausible, general, independently motivatable principles about *oughts*: the subjective ought test and the objective ought transmission principle.²⁵

So I take it that this kind of explanation would be well-motivated and have several nice features, if only we could explain premise 2 in a principled way. In the next section, I'll consider how to do this. Then in part 4 I'll refine the view, and illustrate one of its important advantages. The view developed in parts 1–4 is, however, only roughly correct. In part 5 I'll explain why the assumptions I'm making are only approximately correct, show how to replace them with more defensible assumptions about subjective *reasons*, and show how the resulting view generalizes very neatly to deal with cases of means-end reasoning to non-necessary means—a serious obstacle to many of the existing views in the literature.

3.3 Explaining premise 2

According to premise 2—the missing step in our explanation—if Zach intends to do something, then he subjectively ought to do it. But is this true? And if so, what would it take to explain it? Fortunately, the subjective ought test gives us the answer. According to the left-to-right direction of the subjective ought test, Zach subjectively ought to do *A* only if he has some beliefs which are such that if they are true then Zach objectively ought to do *A*. So if premise 2 is true, then the following thesis about intention *must* also be true:

nature of intention If you intend to do *A*, then you have some beliefs which are such that, if they are true, then you objectively ought to do *A*.

So that tells us what *has* to be true, in order for premise 2 to be true. Moreover, nature of intention is also *sufficient* to explain the truth of premise 2. For premise 2 follows from it, together with the right-to-left direction of the subjective ought test.

So is nature of intention true? The simplest thing that you could believe, in intending to do *A*, that would make nature of intention true, is simply that you ought to do *A*. This, of course, is an old and familiar idea about intentions: that they involve or presuppose judgments about what you ought to do. If this old and familiar idea is true, then that would complete our explanation, by making Zach's case a special case of Xera's:

P1 Zach intends to do *A*.

P2.1 If Zach intends to do *A*, then Zach believes that he ought to do *A*.

P2.2 If Zach's belief that he ought to do *A* is true, then he ought to do *A*.

P2.3 Right-to-left direction of the subjective ought test: if Zach has some beliefs such that, if they are true, then he ought to do *A*, then Zach subjectively ought to do *A*.

P2 If Zach intends to do *A*, then he subjectively ought to do *A*.

²⁵ Alternatively, since the only need we had for the left-to-right direction of the subjective ought test was in order to derive the subjective reflection of the transmission principle from ought transmission, we could make do only with its right-to-left direction, transmission reflection, and premise 2.

P3 Zach believes that to do *A*, he must do *B*.

P4 Transmission Reflection: if Zach subjectively ought to do *A* and believes that to do *A* he must do *B*, then he subjectively ought to do *B*.

C Zach subjectively ought to do *B*.

In this explanation, premises 1 and 3 are simply the features of Zach's situation—in which means-end coherence is supposed to apply. Premise 4 is just the transmission principle for subjective oughts, which is both independently plausible and follows from the transmission principle for objective oughts and the subjective ought test. And premise 2 follows from premises 2.1 to 2.3, of which 2.2 is trivial and 2.3 is the right-to-left direction of the subjective ought test. So in addition to the transmission principle for objective oughts and the subjective ought test, this explanation appeals to only one substantive assumption: premise 2.1, that intending entails believing that you ought. So if this old and familiar assumption is true, then it yields a powerful explanation of what is going on in Zach's case, which otherwise appeals only to very general principles that can be motivated on the basis of independent cases.

4 Defending the idea

4.1 The problem of picking

So far I hope to have argued that the assumption that intending involves believing you ought is a *fruitful* one, in the sense that given the kinds of very general background principles given by the subjective ought test and the transmission principle for objective oughts, it is sufficient to explain why Zach subjectively ought to take the means. There are at least three major problems, however, for the view that intending to do *A* entails believing that you objectively ought to do *A*. In increasing order of difficulty, they are the problem of *picking*, the problem of *akrasia*, and what my colleague Jake Ross calls the *three-envelope* problem. In this section I'll explain the problem of picking and argue that solving it requires no real modification in our view. In the next, I'll explain the problem of *akrasia*, and suggest a few different kinds of modification that would allow us to preserve the outlines of this kind of account of the norm of means-end coherence. I'll put the three-envelope problem off until part 5 and consider a couple of loose ends first, as confronting the three-envelope problem will require revisions to the subjective ought test, as well as to the idea that intending involves believing you ought.

The problem of picking is simple. When Zach goes to buy a carton of milk, there are two for him to choose from, the one on the left, and the one on the right. Neither has any advantage over the other, and since he knows this, he knows full well that it is not the case that he ought to take the one on the left, as opposed to the one on the right. Nevertheless, he has to pick one, and so he opts for the one on the left, *without* believing that he ought to have.

On the face of it, this is an intention. Zach has formed an intention to take the carton on the left. But in truth, it doesn't matter whether we call it an 'intention' or

not. What matters, is that it is subject to the same norms of means-end coherence as the other cases we are trying to investigate. Once Zach opts for the carton on the left in this way, he must, as he recognizes, open the door on the left in order to get it. So something is going wrong with him if he has no intention whatsoever of opening that door! So it actually doesn't matter whether we call this case an intention or not; we won't have solved our problem unless our solution applies in this case. And in this case, it seems that Zach does *not* believe that he objectively ought to take the carton on the left.

I think that what cases of picking really show is that the intention to do *A* needn't require the *antecedent* belief that one ought to do *A*. But I don't think they show that someone who intends to do *A* need not believe that she ought to do *A*. *Before* Zach makes any decision, of course, the carton on the left and the carton on the right are perfectly on a par. But *after* he has made his decision, the carton on the left comes out on top. The fact that he has picked it is now a relevant difference between the two. So if he ought to take one of them, the one on the left is the one that he ought to take.

It is sufficient for this answer that intentions provide reasons—when you intend to do something, that gives you some additional reason to do it, that you did not have before. Call the element of a decision that gives you an additional reason to do it *plumping*. Plumping need only provide you with a very trivial reason, in order to deal with cases of mere picking. So plumping for hiring an assassin to murder your wife is not sufficient to make that what you ought to do. But plumping for the carton on the left *is* sufficient to make that what you ought to do. We do it all of the time. On the view that intention involves the belief that you ought, therefore, Zach is able to intend to take the carton on the left, because he is able to plump for it, and because he understands that his plumping is all that it takes to make a difference between the two. Consequently, though they provide a *prima facie* counterexample, cases of picking present no serious obstacle to the view that intending involves believing that you ought.

4.2 The problem of *akrasia*

A harder problem is that of akratic intentions, which seem not only to be possible, but, as Wallace (2001) and others have emphasized, to be subject to the norm of means-end coherence. In cases of *akrasia*, an agent acts contrary to what she believes she ought to do. So they also appear to be counterexamples to the thesis that intending requires believing that you ought. Exactly how to understand what goes on in cases of *akrasia* is quite a large philosophical problem in its own right; I'll confine myself here to four limited observations.

First, it's worth observing that if it is possible to have contrary beliefs, then one possible view is that someone who intends akratically believes that she ought not to be doing what she intends, but also that she ought to. Her situation might be thought to be analogous to that of the man who deep down believes that his wife is cheating on him, but who, due to wishful thinking, manages to convince himself that she is faithful. If the akratic's situation is like this, then though she suffers from a vice that

is akin to wishful thinking, she is not, after all, a counterexample to the thesis that intending involves believing that you ought. Still, this is not likely to satisfy most, and so I have three observations about how something like the nature of intention can be defended by appeal to weaker assumptions to which the akratic is not a counterexample.

The first way that we could weaken the view, is by weakening what the intender has to believe. Instead of requiring that he believe that he ought to do it, we could instead require that he believe, say, that he has *adequate reason* to do it. Then it would follow, presumably, from an appropriate analogue of the subjective ought test, that intenders have *adequate subjective reason* to do what they intend, and hence by a generalization of the transmission principle, adequate subjective reason to do what they believe to be necessary to what they intend. Or we might weaken the content in some other way. Whichever way, it would yield the same style of explanation, with weaker assumptions.

Another way to weaken the required assumption might be to generalize, and allow that an intender need not *believe* that she ought, so long as she *takes it* that she ought. Many authors have claimed, after all, that intentional action requires that you ‘take yourself’ to have a reason. I’m here assuming that ‘taking’ amounts to a different attitude than belief—perhaps it is quasi-perceptual. Suppose, then, that intention requires *taking* it that you ought. If that is so, then all that we would need would be a small revision in our test for subjective oughts, in order to have the same explanation as before. We have been supposing that your subjective reasons depend on your beliefs. But suppose that we generalized, and held that your subjective reasons depend on what you either believe or *take* to be true. (To motivate this, we might imagine a Ewing-like case of someone who has no *belief* about what he ought to do, but does *take* it that he ought to do *A*.) Given this broader principle, our explanation would go through as before, simply with a weaker assumption about what intention requires. This proposal, I think, may have some promise at addressing the worry about akratic intentions.


Finally, we could drop the assumption that Zach must have a belief about what he ought to do, altogether, and instead require merely that he have some beliefs which are the sort, if they are true, to guarantee that Zach ought to do it. This weaker assumption would guarantee that when Zach intends to do *A*, he subjectively ought to do it, and hence subjectively ought to do what he believes necessary for it, by the subjective reflection of the transmission principle.

All three of these are ways of weakening the crucial assumption that this view about means-end coherence of intention requires. My point here is not to advocate any of them in particular—that would require having a great deal more to say about *akrasia* than I can in the scope of this paper. My point is rather to show that the essential attractions of the explanation in Sect. 3.3 can be preserved under weaker assumptions than premise 2.1. The most I can hope to convince you of here, is that there are grounds for optimism that some combination of these kinds of adjustments can accommodate the problem of *akrasia* in a way that preserves the attractions of this kind of explanation of the norm of means-end coherence, and that the attractiveness of the explanation which results makes it worth at least trying to do so.

4.3 Objection: we've explained the wrong thing

The astute reader will have noticed that I set out to explain why there is something wrong with Zach if he does not *intend* *B*, but what I have actually done, is to explain why there is something wrong with Zach if he does not *do* *B*. His belief that *B* is necessary for *A*, on the view I'm outlining, makes it the case that he subjectively ought to do *B*, due to the transmission reflection and the fact that the intention to do *A* entails the belief that you ought (on the unmodified version of the view). So that explains why he subjectively ought to do *B*, but not why he subjectively ought to *intend* to do *B*.

This is true. I haven't yet explained why Zach subjectively ought to intend to do *B*. But before I say how I would explain this, on the present view, allow me to first draw a picture, in order to contrast my style of explanation with one others have claimed.



Everyone should agree that just as there is an important normative relationship between intending to do *A* and intending to do *B*, there is an important relationship between intending to do *B* and doing it, and an important normative relationship between intending to do *A* and doing *B*. On the usual view, this last relationship, which I have drawn with the diagonal arrow, is compound: it is the result of *composing* the other two relationships, each of which is interesting and distinctive.²⁶ According to this view it is *one* thing to understand the horizontal arrow, and *another* to understand the vertical arrows, and only once we are in a position to understand both, will we be able to put them together in order to understand the diagonal arrow.

Because this view is so common, it can seem surprising that what I have just done, is to have explained the diagonal arrow without first explaining the horizontal one. But having done so, it is natural to reverse the usual order of explanation, and to use our account of the diagonal arrow in order to offer an account of the horizontal one. In fact, doing so is easy. Our explanation so far is this: intending to do *A* makes it the case that you subjectively ought to do *A*. And given this, believing that *B* is necessary for *A* makes it the case that you subjectively ought to do *B*. So by the same principles, believing that intending *B* is necessary for doing *B* makes it the case that you subjectively ought to intend *B*. Just as doing *B* is a means for doing *A*, intending *B* is a means for doing *B*. So from transmission reflection, it follows that if you believe that to do *B* you must intend *B*, then you subjectively ought to intend *B*.

Of course, this does not explain why *anyone* who intends *A* and believes that to do *A* she must do *B* subjectively ought to intend *B*. It only explains this for the special case of people who further believe that to do *B* they must intend *B*. But this

²⁶ For example, see Broome (2001).

is not a bad thing! It is exactly what we need in order to distinguish the cases of *means*, to which the norm of means-end coherence applies, from those of *foreseen side-effects*, which the norm of means-end coherence does not require you to intend.²⁷ If you are the Strategic Bomber and intend to bomb the munitions factory which is next to the school, you believe that the only way that you can bomb the munitions factory is if you bomb the school. Since you subjectively ought to bomb the munitions factory, you subjectively ought to bomb the school. But it is not the case that you subjectively ought to intend to bomb the school, because you don't believe that you need to intend this, in order to bomb it. You think that it will simply happen as a result of bombing the munitions factory.

So, to recap: Zach intends to do *A*, and this entails that he has some belief which, if true, guarantees that he objectively ought to do *A*. So it follows from the right-to-left direction of the subjective ought test that he subjectively ought to do *A*. He also believes that to do *A* he must do *B*, and so from the transmission principle for subjective oughts it follows that he subjectively ought to do *B*. And finally, provided that he believes that to do *B* he must intend *B*, it follows from the transmission principle for subjective oughts that he subjectively ought to intend *B*. And this further assumption about his beliefs is appropriate to make, because once we look closely at cases of unintended side-effects, we see that the norm of means-end coherence doesn't really apply in cases in which he doesn't have this belief, anyway. I'll offer some revisions in part 5, but this is the basic idea behind my proposal for how to account for the norm of means-end coherence on intention.

4.4 A temporal advantage of this account

One advantage of this account is that unlike many accounts of cases like Zach's, it captures the right relationship to his intention over time.²⁸ Whatever goes wrong with Zach when he intends to do *A* but has no intention to do *B*, it is not going wrong with Zach immediately after he decides to do *A* and before he has time to think about what is necessary for doing *A*. It only goes wrong if *at no time* prior to the time at which Zach intends to do *A*, does he form the intention to do *B*. This is something that many accounts of cases like Zach's get wrong. For example, according to Wide-Scope accounts, Zach is at all times under a special requirement to not be such that he both intends to do *A* and does not intend to do *B*. So it follows from such accounts, that whatever is going wrong with Zach is going wrong with him even before he has had a chance to reflect on what is necessary for *A*. So Wide-Scope views seem to predict the wrong results about what happens in Zach's case over time.

The view developed here, however, gets what I think are the *right* results about this case. Since it is only necessary to intend to do *B* at some time prior to the time at which *B* actually needs to be done, there need be nothing strictly wrong with Zach if he has not yet formed the intention to do *B*. Though he subjectively ought to form

²⁷ I take this moral to be familiar from the work of Michael Bratman, who has pressed it particularly acutely.

²⁸ Compare Raz (2005b, p. 5).

this intention, he is not *strictly* required to do it right away. Since that seems like the right result, and other views seem to get the wrong result, or to build in the right result only through stipulation, I think that is collateral evidence that this view about the norms of means-end coherence on intention is worth taking seriously.

5 Generalizing

5.1 Reconsidering the subjective ought test

I think I have illustrated enough for us to see the appeal of this view. It gives us an answer to the question of *in what sense* there is something going wrong in Zach's case when he does not intend *B*—a problem that the failures of Wide-Scope accounts and of the naïve view show is hard. Moreover, the answer that it gives us to the *in what sense* question also leads to an answer to the *why* question that appeals, with one exception, only to plausible independent principles: that there is a distinction between objective and subjective 'oughts' which are related in something roughly like the way indicated by the subjective ought test, and that objective oughts transmit to necessary means. The only other necessary assumption is that an agent who intends to do *A* necessarily satisfies some condition which would give her a subjective reason to do *A*, and in part 4 I suggested that some version of this claim might very well be defensible, though I have not had the space, here, to go into defending some version of this hypothesis in detail. Certainly it is the sort of thing that philosophers have believed for independent reasons.²⁹

I've formulated the account so far by appeal to the subjective ought test as formulated in Sect. 2.1. This has been because that principle has been accepted by many philosophers, including philosophers with different views about, for example, the priority between the subjective and objective 'ought' relations, and because it is easy to motivate by consideration of initial examples like that of Bernie. But in fact, however, I believe that the subjective ought test is correct only in spirit. I'll now explain why I do not believe that it can actually be correct. Then in Sect. 5.2 I'll briefly defend what I think is right, instead, and use that to show that the essentials of the account of means-end coherence are unaffected. Finally, in Sect. 5.3 I'll argue that the modified account has an advantage not shared by most existing accounts of the norm of means-end coherence on intention.

According to the subjective ought test formulated in Sect. 2.1, if Yves believes that he objectively ought to do *A*, then Yves subjectively ought to do *A*. This is because, in believing that he objectively ought to do *A*, he believes something which guarantees (trivially) that he objectively ought to do *A*. From the same test, it also follows that Bernie subjectively ought to take a sip, conditional on our assumptions that Bernie believes that his glass contains gin and tonic, and that if it contains gin and tonic, then he objectively ought to take a sip. But now suppose that Yves *is* Bernie and that *A is* the action of not taking a sip. That is, imagine that Bernie believes that his glass contains gin and tonic, but also, for some reason, believes that

²⁹ Compare, for example, Davidson (1978) and Tenenbaum (2007).

he ought not to take a sip. Our test as so far formulated predicts that Bernie subjectively ought to take a sip, and also that he subjectively ought to not take a sip.

This seems bad. It should not turn out that Bernie subjectively ought to do both things. Intuitively, the problem is that what Bernie subjectively *ought* to do should turn on *all* of his beliefs, not simply on one or two. What he subjectively *ought* to do should turn on what happens when you put the effects of all of his beliefs together.

But on the other hand, tests for the relationship between objective and subjective ‘oughts’ that turn on the whole sum of an agent’s beliefs are subject to other kinds of counterexamples. For example, Jacob Ross considers the following sort of case³⁰: Wynn has the opportunity to choose one of three envelopes set in front of her. Whatever it contains, she will be able to keep, and she will not get what is in the other two envelopes. She believes that the first envelope contains 200 dollars, that one of the other envelopes contains three hundred dollars, and that the other contains nothing. And she considers it equally likely that the 300 dollars are in the second envelope as in the third. According to Ross, she subjectively ought to take the first envelope, because it has the highest expected value, given her beliefs. But if her beliefs are all true, then one of the following has to be the case: either there is more money in the second envelope than in the first, or there is more money in the third envelope than in the first. Either way, it would be a mistake to take the first envelope. The thing she *objectively* ought to do, is the one that will actually be best. It follows that it can’t be true that Wynn subjectively ought to do whatever it would be the case that she objectively ought to do, were all of her beliefs true. *Incompleteness* in her beliefs can play a role in fixing what she subjectively ought to do.


As I anticipated in part 4, the three envelope problem is also a serious obstacle to the thesis that intending to do *A* requires believing that one objectively ought to do *A*. When Wynn chooses an envelope, she does so despite knowing full well that she objectively ought to choose one of the others—she simply doesn’t know which.

Finally, I’m independently wary about the tests which posit a direct relationship between the objective and subjective ‘oughts’. This is because I believe not only that there are objective and subjective ‘oughts’, but that there are objective and subjective reasons. And it seems to me that objective oughts are related to objective reasons in the same way that subjective oughts are related to subjective reasons. Finally, I hold that oughts are to be accounted for in terms of reasons, rather than conversely.³¹ You ought to do what the balance of your reasons favors, I hold. It follows from these views that there can’t be a direct reductive relationship between the objective and subjective ‘oughts’ that would validate anything like our tests. If objective oughts are to be accounted for in terms of objective reasons, and subjective oughts are to be accounted for in terms of subjective reasons, then neither objective oughts nor subjective oughts can directly be accounted for in terms of the

³⁰ Ross (2006). Broome (unpublished) and Parfit (forthcoming) also discuss such a case. The original version of such a case is attributed to a footnote in Regan (1980).

³¹ The converse view is articulated by, for example, Toulmin (1950) and Broome (2004).

other. The following picture shows how I think things work (the arrows, read from head to tail, represent ‘is analyzed in terms of’):



I can't, of course, defend each of these views here, but I think they are relevant to a complete appreciation of the grounds for thinking that objective oughts and subjective oughts might be directly related.

5.2 The correct view, framed in terms of reasons

All of these problems are solved, I think, by the idea that it is objective and subjective *reasons* which are directly related, and related by something very much like the subjective reason test formulated in Sect. 2.3. On this view, though it does not follow immediately from the fact that Bernie believes that his glass contains gin and tonic that he ought to take a sip (for he might also believe that he ought not), it does follow that he has a subjective *reason* to take a sip. That is why, if other things are equal, it *will* be the case that he subjectively ought to take a sip. For what he subjectively ought to do, is simply whatever he has most subjective reason to do. So this solves our first problem.

It also addresses Ross's kinds of case. In Ross's case, if Wynn's beliefs are all true, then it is not the case that she objectively ought to take the first envelope. It is either the case that she objectively ought to take the second, or that she objectively ought to take the third. On my view, this is because what she objectively ought to do is a matter of *all* of her objective reasons, which include the following: the fact that there is \$200 in the first envelope is a reason to take it, and the fact that there is \$300 in the second envelope is a reason to take it. But since the second reason is a weightier reason than the first, and there are no other relevant reasons, she objectively ought to take the second envelope. The same thing goes, if the \$300 is in the third envelope instead of the second.

But what Wynn subjectively ought to do, on my account, is a matter of the weight of her subjective reasons. And those include the following: her belief that the first envelope contains \$200 gives her a reason for her to take the first envelope, her belief that the second envelope is 50% likely to contain \$300 gives her a reason for her to take the second envelope, and her belief that the third envelope is 50% likely to contain \$300 is a reason for her to take the third envelope. But since she has no view about which of the second or third envelope contains the \$300, she has no further belief to figure as a reason. Now which of these subjective reasons is weightiest? Well, suppose they are all true. The fact that the first envelope contains \$200 would seem to me to be a weightier reason for her to take it, than the fact that the second envelope is 50% likely to contain \$300. So just based on *those* two reasons, the thing for her to do would be to take the first envelope. Assuming that

subjective reasons are as weighty as their corresponding objective reasons would be, therefore, it follows that Wynn's subjective reason to take the first envelope is the weightiest.³²

Of course, if these beliefs are true, then some further fact must also be true, to the effect that the second envelope either does, or does not, contain the \$300. But since Wynn does not have a belief about that, it does not figure among her subjective reasons by way of the converse reason test. So I conclude that what Wynn subjectively ought to do, as Ross insists, is to take the first envelope. We actually predict this result, if we take the direct connection to hold between objective and subjective reasons, rather than between objective and subjective oughts.

Moreover, switching from oughts to reasons dissolves the problem raised by the three envelope problem for the thesis that intending involves believing you ought. True, due to the three envelope problem, it is not tenable to claim that intending to do *A* requires believing that you objectively ought to do *A*. But it may still require believing, or taking it that, you have a sufficiently weighty objective reason to do *A*.

Finally, these revisions still suffice for the account of Zach's case, basically substituting 'reason' everywhere in our earlier explanation for 'ought'. On the assumption that intending to do *A* entails believing that you have a reason to, it follows from the right-to-left direction of the subjective reason test that if you intend to do *A*, then there is a subjective reason for you to do *A*. What we then need is a subjective reason transmission principle, which I suggest that we can derive from the subjective reason test together with the generalization of the objective ought transmission principle to the case of reasons. Here I'll state what I take to be a special case:

reason transmission If *X* has an objective reason to do *A* and to do *A* *X* must do *B*, then *X* has an objective reason to do *B* of equal weight to *X*'s objective reason to do *A*.

In the same way as for the transmission principle for objective oughts, we can use this to derive its subjective reflection. (Actually, the derivation requires principles which include reference to weights, which is stronger than the principles stated in Sect. 2.2.) I'll just state the transmission principle that I think is plausible:

reason trans reflect If *X* has a subjective reason to do *A* and *X* believes that to do *A* she must do *B*, then *X* has a subjective reason to do *B* of weight at least as great as *X*'s subjective reason to do *A*.

³² Officially, my views about the weights of reasons are more complicated than this. See Schroeder (2007b, Chap. 7). But I don't think that I am committed to anything which would make this principle fail in the simple cases that we are considering. A helpful referee also encourages me to note that determining the weights of subjective reasons, as I do in an informal and intuitive way throughout this section and the next, is going to be a complicated matter—particularly because they can derive from sets of beliefs which are not themselves consistent. There is unfortunately insufficient space to take this issue up here in sufficient depth, here, so I'll reserve it for a future occasion. My goal in this paper has been to stress the attractions of this basic idea, and there remains considerable work to be done both in defending the required thesis about the nature of intention, and in articulating and precisifying the intuitive claims I'm making about the weight of reasons, both objective and subjective.

It follows from this principle that Zach has a subjective reason to do *B* that is at least as weighty as whatever subjective reason he has on account of intending to do *A*. This accounts, I claim, for the *stringency* of the norm of means-end coherence, the feature that I claimed in Sect. 2.2 was central to what is difficult about this case. Means-end coherence is *stringent* because when you believe that the means is *necessary* for the end, the full force of your subjective reason to pursue the end transmits to your subjective reason to pursue the means.

5.3 A final advantage: generalizing from necessary means

Finally, I want to show how this allows us to bring out a final important feature of this account. Just as there is some interesting relationship between intending ends and intending means which are necessary to those ends, so also there is a more general connection between intending ends and intending non-necessary means that nevertheless *facilitate* those ends—which help to bring them about. But most existing views about the norm of means-end coherence on intention trade essentially on features of the paradigm case, in which the agent's means-end belief is that the means are *necessary* for the end.

For example, on Harman's (1976) account, on which many other accounts have since been modeled, you ought to intend the means you believe are necessary, because the intention to do *A* entails the belief that you will do *A*. And if you believe that you will do *A* and that if you do *A* you will do *B*, then theoretical rationality requires the belief that you will do *B*, and it is because you are required to believe this, that you are required to intend the means (the further details aren't essential for the point that I want to make). This models the means-end relation on *modus ponens*. But that seems to require a *deductive* relationship between end and means. It makes it hard to see how it will generalize to the more general case in which Zach has beliefs about what would facilitate his end which he doesn't take to be strictly necessary for it. Similar features are shared by the accounts offered by Wallace (2001), Broome (2001), and Setiya (2007b).

On the account offered here, on the other hand, it is easy to see why things should generalize, and why we will get the intuitively right results when we generalize. This is because the transmission principle for reasons that I have stated is really only a special case. In general:

general reason transmit If *X* has an objective reason to do *A* and *X*'s doing *B* would facilitate her doing *A*, then *X* has an objective reason to do *B* of weight at least proportional to *X*'s objective reason to do *A*, and to how well her doing *B* would facilitate her doing *A*.³³

And this, of course, yields, together with suitable versions of our tests for subjective reasons and their weights:

³³ Compare Raz (2005a). I take these this principle to be intuitive, even though I have no way of making the 'proportionality' claim precise.

subjective version If X has a subjective reason to do A and believes that her doing B would facilitate her doing A , then X has a subjective reason to do B of weight at least proportional to X 's subjective reason to do A , and to how well she believes her doing B would facilitate her doing A .

But by now it should be obvious that we can use such a principle in order to explain subjective reasons for Zach to take what he believes to be the facilitative means to his intended ends. Moreover it predicts, correctly, that these subjective reasons will not be as weighty as the subjective reasons to take the means he believes to be necessary. And so it allows that this case does not exhibit the stringency of the necessary means case. All of these predictions are intuitively correct.

In my view, it should be a constraint on an adequate account of Zach's case—of the so-called norm of means-end coherence—that it both explain the special stringency of what is going on when Zach's belief is that some means is necessary for his end, and also generalize to explain what is going on in cases where he has beliefs in non-necessary but facilitative means. So I take it to be an important and distinctive virtue of this account that it looks like it can do this correctly, and by appeal to relatively weak assumptions.

To reiterate: what we need for this account is some general relationship between objective and subjective reasons along the lines of the subjective reason test formulated in Sect. 2.2, along with corresponding claims about those reasons' weights. Such a principle is motivatable entirely independently of the problem about means-end coherence, and I argued in Sects. 5.1 and 5.2 that its correct version withstands problems for less careful formulations. We also need to appeal to something in the neighborhood of general reason transmit, which looks like a highly plausible and totally general principle, at least something like which we should independently accept. And finally, we need to accept some hypothesis about the nature of intention—for example, that intending to do something requires taking yourself to have a reason to do it.

I haven't argued for this final hypothesis. But I'm definitely convinced that the considerations adduced here show that it would be an explanatorily powerful assumption, if it were true. I know of no other account either of *in what sense* there is something wrong going on in Zach's case, nor of *why* this is so, which is entirely happy. And this account not only avoids the pitfalls for others, but deals correctly with the way that the norm of means-end coherence applies over time, and generalizes in the right kind of way to accommodate merely facilitating means. So the hypothesis seems well worth exploring further than I have been able to do, here.³⁴

³⁴ The central ideas of this paper were formulated during Michael Bratman's presentation to the 2006 conference on Practical Reason at Bowling Green State University. My apologies to him for being distracted during the talk. I also owe thanks to Bratman, Doug Lavin, and Niko Kolodny for listening to my incoherent formulations that weekend, to Joseph Raz, to Kieran Setiya's enlightening paper, to conversations with Jacob Ross, to very helpful comments from Jamie Dreier, Steve Finlay, and Jacob Ross, to an excellent and insightful referee, and to an audience at Georgetown University.

References

- Bratman, M. (1981). Intention and means-end reasoning. *Philosophical Review*, 90(2), 252–265.
- Bratman, M. (1987). *Intentions, plans, and practical reason*. Cambridge: Harvard University Press.
- Bratman, M. Intention, belief, theoretical, practical. In J. Timmerman, J. Skorupski, & S. Robertson (Eds.), *Spheres of reason* (forthcoming).
- Broome, J. (1999). Normative requirements. *Ratio*, 12(4), 398–419.
- Broome, J. (2001). Normative practical reasoning. *Proceedings of the Aristotelian Society*, 75, 175–193.
- Broome, J. (2004). Reasons. In Wallace, R. J., et al. (Eds.), *Reason and value: Themes from the moral philosophy of Joseph Raz* (pp. 28–55). Oxford: Oxford University Press.
- Darwall, S. (1983). *Impartial reason*. Ithaca: Cornell University Press.
- Davidson, D. (1978). Intending. Reprinted in his *Essays on actions and events* (pp. 83–102). Oxford: Oxford University Press.
- Ewing, A. C. (1953). *Ethics*. London: English Universities Press.
- Gensler, H. (1985). Ethical consistency principles. *Philosophical Quarterly*, 35(3), 156–170.
- Greenspan, P. (1975). Conditional oughts and hypothetical imperatives. *The Journal of Philosophy*, 72(10), 259–276.
- Harman, G. (1976). Practical reasoning. Reprinted in his *Reasoning, meaning, and mind* (pp. 46–74). Oxford: Oxford University Press, 1999.
- Hill, T. (1973). The hypothetical imperative? *The Philosophical Review*, 82(4), 429–450.
- Kolodny, N. (2005). Why be rational? *Mind*, 114(3), 509–563.
- Kolodny, N. (2007). State or process requirements? *Mind*, 116(2), 371–385.
- Parfit, D. *Climbing the mountain*. Oxford University Press (forthcoming).
- Raz, J. (2005a). The myth of instrumental rationality. *Journal of Ethics and Social Philosophy*, 1(1), 2–28, <http://www.jesp.org>
- Raz, J. (2005b). Instrumental rationality: A reprise. *Journal of Ethics and Social Philosophy*, 1, 2–20, symposium, <http://www.jesp.org>
- Regan, D. (1980). *Utilitarianism and cooperation*. Oxford: Oxford University Press.
- Ross, J. (2006). *Acceptance and practical reason*. PhD dissertation, Rutgers University.
- Schroeder, M. (2004). The scope of instrumental reason. *Philosophical Perspectives*, 18(Ethics), 337–364.
- Schroeder, M. (2005). Instrumental mythology. *Journal of Ethics and Social Philosophy*, 1, 2–12, symposium <http://www.jesp.org>
- Schroeder, M. (2007a). Weighting for a plausible humean theory of reasons. *Noûs*, 41(1), 138–160.
- Schroeder, M. (2007b). *Slaves of the passions*. Oxford: Oxford University Press.
- Schroeder, M. Having reasons. *Philosophical Studies* (forthcoming).
- Setiya, K. (2007a). *Reasons without rationalism*. Princeton: Princeton University Press.
- Setiya, K. (2007b). Cognitivism about instrumental reason. *Ethics*, 117(4), 649–673.
- Tenenbaum, S. (2007). *Appearances of the good: An essay on the nature of practical reason*. Cambridge: Cambridge University Press.
- Toulmin, S. (1950). *Reason in ethics*. Cambridge: Cambridge University Press.
- Wallace, J. (2001). Normativity, commitment, and instrumental reason. *Philosophers' Imprint*, 1(3), <http://www.philosophersimprint.org/001003>
- Williams, B. (1981). Internal and external reasons. Reprinted in his *Moral Luck* (pp. 100–110). Cambridge: Cambridge University Press, 1981.