

*Do Intentions Change Our Reasons?**

Niko Kolodny

Attitudes matter, but in what way? How does having a belief or intention affect what we should believe or intend? One answer is that attitudes themselves *justify*, or give us *reason* for, other attitudes. Like many others, I find this answer, at least on reflection, puzzling.¹ Our beliefs are justified by considerations that suggest that they are true, and our intentions are justified by considerations to the effect that they might lead us to act well. Why should having one belief, in itself, generally suggest that another belief is true, or having one intention, in itself, generally mean that another intention would achieve something worthwhile? If attitudes matter, they seem to matter in a different way. If we have one attitude, then we are *internally incoherent* insofar as we fail to have other attitudes. In this way, attitudes may make us criticizable as *irrational* for failing to have other attitudes. But attitudes do not, themselves, justify, or give us reason for, other attitudes. Or so “nonpsychologism,” to give it a name, claims.²

* I am very grateful to private comments from Tim Scanlon, Sam Scheffler, Jay Wallace, and Wai-Hung Wong, as well as to public comments from participants in the Seminar on Ethics and Normative Theory (SENT) at Stanford and in a seminar at the University of Valencia, where Josep Corbi and Eduardo Ortiz gave particularly insightful prepared remarks.

¹ John Broome, “Are Intentions Reasons? And How Should We Cope with Incommensurable Values?” in Christopher Morris and Arthur Ripstein, ed., *Practical Rationality and Preference: Essays for David Gauthier* (Cambridge: Cambridge University Press, 2001), pp. 98–120; Johnathan Dancy, *Practical Reality* (Oxford: Oxford University Press, 2000); Derek Parfit, *Climbing the Mountain* (unpublished); Joseph Raz, “The Myth of Instrumental Rationality,” *Journal of Ethics and Social Philosophy* 1:1 (2005); T.M. Scanlon, “Reasons: A Puzzling Duality?” in R. Jay Wallace, Philip Pettit, Samuel Scheffler, and Michael Smith, ed., *Reasons and Value: Themes from the Moral Philosophy of Joseph Raz* (Oxford: Clarendon Press, 2004), pp. 231–246; and R. Jay Wallace, “Normativity, Commitment, and Instrumental Reason,” *Philosophers’ Imprint* 1:3 (2001). Michael Bratman, *Intention, Plans, and Practical Reason* (Cambridge: Harvard, 1987); and John Brunero, “Are Intentions Reasons?” (draft) agree that intentions, at least, do not give us reasons, which is the topic of this paper.

² It is difficult to argue for nonpsychologism in a way that goes beyond the point that having such attitudes seems, on reflection, not to carry any justificatory weight. Consider the claim that will be of interest to us: that intending E gives one reason to take means to E-ing. A certain kind

Nonpsychologism provokes a variety of doubts. Some of these doubts are broadly metaethical. What could it be for some consideration to be a reason for us, if not that we had some psychological investment in it? How else could reasons explain, in the way they do, our psychological responses to them? Although such doubts are clearly important, I will be concerned with a different kind of doubt, which, in a way, arises earlier. Can nonpsychologism make sense of our ordinary thought? Setting aside whatever theoretical, second-order mysteries it may court, is nonpsychologism even viable as an interpretation of our first-order normative consciousness?

Consider the familiar phenomenon that deciding to do something affects what steps one ought to take going forward. Suppose, for example, that I have merely sufficient reason to attend

of “bootstrapping” argument begins by describing a situation as inhospitable as can be to one’s having reason take the means to E-ing: a situation in which E-ing, and taking the means to E-ing, is utterly pointless or appalling. It then adds to the description of the situation the fact that one intends to E and invites us to agree that we still lack reason to take means to E. An ambiguity between sufficient and *pro tanto* reason, however, makes this less decisive than it might at first seem. While it may be clear that one still lacks sufficient reason to take the means to E, it is less clear that one has not acquired at least a *pro tanto* reason to take the means to E, which is outweighed by the pointless or appalling character of the means. We might, in reply, be asked to consider a case in which the reasons to take the means to E and not to take the means to E are otherwise perfectly balanced. If intending to E gives one a *pro tanto* reason to take the means to E, then it would, in the context of the otherwise perfect balance, give one conclusive reason to take the means to E. The problem is that this is not obviously implausible. Indeed, I will be trying to explain a phenomenon that resembles it. A different kind of “bootstrapping” argument aims to show that psychologism implies that attitudes are reasons for themselves. Yet the claim in question implies not that intending E gives one reason to intend to E, but instead that intending to E gives one reason to take the means to E. (It is true that, typically, *sustaining* an intention to E is a necessary means to E-ing. So although intending at t to E would not give one reason to intend at t to E, it would give one reason to continue to intend *going forward* from t to E. But this is not to say that intending at t to E gives one reason to intend *at t* to E.) At any rate, it is unclear why someone who does not already reject the view that intending E is a reason for *other* attitudes should be troubled by the implication that intending E is a reason for *this* attitude: intending to E. It might be suggested that if intending to E gives one reason to intend to E, then one cannot but conform to this reason. But it is not clear why it must be possible to violate a reason. The thought, perhaps, is that once we agree that given that p, one cannot but X, it becomes somehow pointless to tell one that p is a reason for one to X. But this seems more a point of pragmatics than semantics.

a conference in New York, merely sufficient reason to attend a conference in Boston, and conclusive reason to attend exactly one of these. As things stand, I haven't made up my mind whether to go. The New York organizers call me to ask whether I would like to pre-register, which, as they explain, is cheaper than registering on site. It would be reasonable for me to respond: "Sorry, I don't know whether I'm going, so I don't have any reason to register yet." Suppose that I go on to make up my mind to go, and the organizers call back. Now, it seems, I do have sufficient reason to pre-register for New York, whereas I still lack sufficient reason to pre-register for Boston. Given that I have decided to attend New York, the fact that pre-registering for New York would serve, in an efficient way, attending New York is a stronger reason than it was before.

If we lack sufficient reason for the end in the first place, it seems right to say, then intending it does not give us greater reason to take the means to it. If an end is pointless or appalling, then intending it does nothing to strengthen the case for taking the means. And if we have conclusive reason for the end, then it is unclear whether intending the end gives us greater reason to intend the means. If it is the only thing to be done, then perhaps we already have all of the reason we could ever need to take the means to it. However, when we have sufficient reason, but only sufficient reason, for the end, it seems right to say that once we have adopted that end, we come to have reason to take the means to it that we did not have before. What, after all, is the alternative? Am I instead to think that I *continue to have just as much* reason to take steps to attending the conference that I have decided *against* attending? Should I just as soon book a flight to that conference, as to the conference that I do plan to attend? To a first approximation, then, the phenomenon is that of:

Intention Effect, first pass (IE1): If I have merely sufficient reason to E, then if I intend to E, I have stronger reason than I otherwise would to take means to E-ing.³

The question that I pursue in this paper is whether we can account for IE1 within the confines of nonpsychologism: that is, without appealing to the idea that, at least in conditions such as these, having an intention *itself* gives us stronger reason to take means to it.

To be sure, we need not accept that intentions themselves are reasons in order to accept that *worthwhile ends* provide reasons for means. They do, in ways that Joseph Raz summarizes in his:

Facilitative Principle (FP): If I have reason to E, then I have reason to take some sufficient means to E.⁴

But this is not the phenomenon that IE1 describes. As far as FP is concerned, intending E makes no difference. What gives me reason to take the means is the value of the end, not my intending it. What IE1 reflects, by contrast, is that intending the end does affect my reason to take the means.

Nor need we accept that intentions themselves are reasons in order to accept that intending an end makes it *irrational* of one to fail to have other attitudes toward the means. For example, many who deny that intentions themselves are reasons accept:

Means-End (ME): It is irrational of me (to intend at t to E, to believe at t that intending at t to M is a necessary means to my E-ing, but not to intend at t to M).⁵

³ Derek Parfit, “Motives,” unpublished; Raz (p.22–23); and Scanlon (p. 236–7) describe this phenomenon. Parfit writes: “There are often many possible aims that would be equally good... In such cases, we cannot have, or pursue, all these aims. Since that is so, what we have most reason to do may depend on which, among these possible aims, we have made *our* aims.”

⁴ “The Myth of Instrumental Rationality,” XX. Raz’s full formulation includes some further qualifications.

But, again, this is not the phenomenon that IE1 describes. IE1 has to do with our *reasons*, rather than our *rationality*.

However, it might be suggested that we are inclined to accept IE1 only because we confuse it with some principle of rationality.⁶ To begin with, I doubt that this principle could be ME. ME applies only to what I believe are *necessary* means, whereas in this case the means are not necessary. Pre-registration is not required; it is just cheaper than the alternative. Next, ME requires me only to *intend* the means. However, what seems out of place in the case described is not simply to refuse to intend the means, but also to refuse to *see* myself as having greater *reason* to *take* the means. Suppose that I have decided to attend the New York conference. Whether or not I intend to pre-register for New York, it would be odd of me to maintain that, so long as I am decided on New York and against Boston, it is a matter of indifference whether I pre-register for New York or for Boston instead. Finally, ME applies to me even when I believe that I lack sufficient reason for the end. IE1, by contrast, is predicated on my having sufficient reason for the end.

If we were confusing IE1 with a principle of rationality, it seems, that principle would have to be:

⁵ Bratman, *Intentions, Plans, and Practical Reason*, XXX; Broome, “Are Intentions Reasons?” XXX; Wallace, “Normativity, Commitment and Instrumental Reason,” XXX. At times, Bratman speaks of intentions giving the agent “framework reasons,” but I take these “framework reasons” to be equivalent to rational requirements. At any rate, these framework reasons are sufficiently different from the kinds of reasons involved in IE1 that they cannot explain the phenomenon. For one thing, these framework reasons seem to be relative to the agent’s means-end beliefs in a way in which the reasons involved in IE1 are not. See the following paragraphs.

⁶Broome, “Are Intentions Reasons?” XXX and Scanlon, “Reasons: A Puzzling Duality?” XXX suggest this.

Believed Intention Effect (BIE): If I believe that I have merely sufficient reason to E and intend to E and believe that M-ing is a means to E, then it is irrational of me to deny that my intending to E gives me stronger reason than I would otherwise have to M.

The problem is not that there is no such principle of rationality. There may well be. The problem is instead that BIE seems to *presuppose* IE1, read literally, as a distinct fact about reasons. Suppose that IE1 were false. Now consider a reflective agent who *knew* that IE1 was false: that after deciding to E, one does *not* come to have stronger reason. According to BIE, rationality would require her knowledge: to believe, falsely, that in deciding to E, she comes to have reason that she lacked before. How can rationality, as a rule, require us to take a false view of our reasons?⁷

There is a more basic problem with the suggestion that ME, or BIE, or any other principle of rationality might explain the phenomena that lead us to accept IE1. These principles require something of the agent when (provided the other psychological conditions are met) and only when he *believes* that M (or intending to M) is a means to E. Since an agent's rationality consists merely in the coherence of his attitudes, what matters for his rationality is merely what he believes, not what is the case. By contrast, IE1 does *not* apply when the agent *falsely* believes that M-ing is a means to E-ing. If in fact there is no pre-registration for the conference, and the agent has simply misunderstood the situation, then we spectators, realizing this, will *not* think that the agent's decision makes it the case that he has reason to pay. We might advise him of this. And IE1 *does* apply even when the agent *fails to realize* that M-ing is a means to E-ing. If

⁷ Scanlon, "Reasons: A Puzzling Duality?" suggests such a principle of rationality (p. 236). But he does not suggest that IE1 is plausible only because it is confused with this principle. He agrees that this principle of rationality must be backed up by a truth about reasons, for much the same reasons as given in the text.

the agent does not realize that he must register for the conference,⁸ then we will *still* think that his decision makes it the case that he has reason to pre-register. Again, we might advise him of this. IE1 appears to be a claim about reasons, rather than a claim about internal, psychic harmony.

How, then, can a nonpsychologist explain IE1? Let us begin with an analogy to the case of belief. Nonpsychologists hold that the fact that I have some belief is not itself a reason to have some other belief. Nevertheless, nonpsychologists presumably do not mean to deny that, say, the fact that I have the paranoid belief that everyone is out to get me is, say, (1) a reason for me to believe that my anti-psychotic medication is wearing off, or (2) a reason for me to believe that someone believes that everyone is out to get him. What they mean to deny is instead that my having that belief is, say, (3) a reason for me to believe that my doctor is out to get me. It does not justify that belief. It means, at most, that if I refuse to believe it, then I am being irrationally incoherent.

What distinguishes (1) and (2) from (3)? The fact that I believe that everyone is out to get me is a reason for me to believe that my medication has worn off just as the obtaining of any event that usually has a certain cause is a reason to believe that that cause obtains. Likewise, the fact that I believe that everyone is out to get me is a reason for me to believe that someone believes that everyone is out to get him, just as any fact that a particular has a property is a reason for me to believe the corresponding existential generalization. My belief is a reason, one wants to say, but *as something else*, not *as my belief*. What nonpsychologism denies is that my attitudes are reasons *as my attitudes*.

⁸At least if this is something that, epistemically speaking, he ought to realize. See section 4.

Yet what is it for my attitude to affect my reasons, but not *as* my attitude?⁹ Here is one interpretation (which is not meant to preclude others): that my having an attitude affects my

⁹ Scanlon appears to recommend something along these lines. He stresses that adopting a goal does not “generate” or “create” a reason to take means to it. Instead, it puts one in a “special relation” to the goal, just as, say, developing specialized skills to pursue that goal would put one in a special relation to it. It is in virtue of this special relation that one has stronger reason. He does not explain what this special relation is supposed to be, however, and it is not clear that he intends it to fit the interpretation that I go on to suggest. That is, it is not clear that he believes that the kind of relation to a goal that adopting a goal puts one in is or must be a kind of relation that a fact that did not consist in one’s adopting that goal (or having any other attitude) might also put one in.

Raz suggests that adopting a goal can affect one’s reasons in a way that *only* adopting a goal can affect one’s reasons. Adopting a worthwhile goal can “activate conditional reasons that we have anyway” to perform actions “whose point depends on being taken as part of the pursuit of some goal” (p. 23). It is not clear how this is compatible with nonpsychologism, if it is even meant to be. It is elusive what difference there might be between the claim that adopting a worthwhile goal (but only a worthwhile goal) generates or creates further reason and the claim that it “activates... reasons that we have anyway” that are conditional on precisely our adopting that goal. The difference cannot be, as “anyway” seems to imply, that on the latter option, but not the former, we have the reasons *before* adopting the goal. We don’t have a conditional reason, strictly speaking, until the condition is met. Of course, it is true, before adopting the goal, that if we meet the condition, then we will have such a reason. But the same seems to be true on the former, more “psychologist” option.

Raz gives the following illustration of the phenomenon: “For someone intent on running a marathon every day during August, running a marathon today, the 20th of August, is crucial to the realization of his ambition. For me it is just an opportunity to know what running a marathon feels like — a matter of much less moment.” The case is complicated, however, by factors that are independent of the adopted goals. If I haven’t run a marathon on each of the first 19 days of August, then, whatever my goals are, my running a marathon today does not bring me any closer to achieving the end of running a marathon every day in August. If the August marathoner has run a marathon on each of the first 19 days of August and remembers what running them felt like, then, whatever his goals, his running a marathon today does not bring him any closer to achieving the end, long since achieved, of acquiring knowledge of what running a marathon feels like. So let us suppose that the date is August 1, neither of us knows what running a marathon feels like, and yet each of us is capable, and equally capable, of running a marathon every day during August. Why, then, does the fact that I have only the goal of knowing what a marathon feels like, whereas he has the goal of running a marathon every day during August, mean that he has stronger reason than I have to run a marathon today? Raz does not elaborate.

Here is a possible answer, along the lines that I go on to develop. Given that I have no plans to run a marathon every day in August, I almost certainly won’t run a marathon every day in August, even if I run one today. And if I won’t run a marathon every day in August, then running a marathon today will not bring me any closer to achieving that goal. As things are, it is no means at all to that end. My running a marathon today is instead a sufficient, but not

reasons in the same way in which some other fact, that did not consist in my having some attitude, might affect my reasons. Somewhat more precisely: It is compatible with nonpsychologism that I have more (or less) reason for attitude B because I have attitude A, if there is some principle of the form:

If some proposition, p , of type P , is true, then, because of this, there is more (or less) reason for me to give some response $R(p)$

where R is a function from propositions to responses, such that

- (i) that I have attitude A is, or makes true, some p of type P ,
- (ii) $R(\text{that I have attitude A})=B$,

and, crucially, for nonpsychologism

- (iii) there is some proposition, p^* , of P such that for no attitude A^* , p^* is that I have A^* .

The principle relevant to (1) is something like:

Abduction: If a usual effect of some event occurs, then because of this, there is reason for me to believe that this event has occurred.

And the principle relevant to (2) is something like:

Generalization: If a particular instantiates some property, then because of this, there is reason for me to believe the corresponding existential generalization.

Propositions other than that I have some attitude can satisfy the antecedents of Abduction and Generalization. So (iii) is satisfied. By contrast, in case (3), the relevant principle would seem to be:

necessary means to the less momentous achievement of knowing what running a marathon feels like. By contrast, given that the August marathoner does have plans to run a marathon every day in August, there is a chance that he will do it. His running a marathon today is thus a necessary means to the more momentous achievement of running a marathon every day of August.

[Another possibility: A history of investment in a project like a history of investment in a friendship.]

If I believe something, then, because of this, there is reason for me to believe its obvious consequences.

(iii) would not be satisfied, since the only facts that can satisfy the antecedent are facts that I have some attitude: namely, that I believe something. According to nonpsychologism, therefore, there is no such principle. The closest principle, if there is one, is of the form:

If I believe something, then, because of this, it is irrationally incoherent of me to refuse to believe one of its obvious consequences.

What the nonpsychologist needs to do in order to account for IE1, therefore, is to specify a principle of the form:

If some proposition, p , of type, P , is true, then, because of this, there is more reason for me to respond in some way $R(p)$,

where:

- (i) that I intend to E is, or makes true, some p of P
- (ii) $R(p) = \text{to } M$
- (iii) there is some proposition, p^* , of P such that for no attitude A , p^* is that I have A .

It remains to be seen what this principle might be.

2. It is often noted that if one has merely sufficient reason to E , *taking* some means to E can give one greater reason, on balance, to E going forward. Having already taken some means toward the end, one can realize it, going forward, with less cost. So, on balance—that is, taking into account the costs and benefits—there is greater reason to E .¹⁰ Since this mechanism is predicated on taking means to E , rather than on intending to E , however, it might seem, at first

¹⁰ See, for example, Broome, “Are Intentions Reasons?” pp.???, Raz, “The Myth of Instrumental Rationality,” p. 22.

glance, irrelevant to IE1. But this may be too quick. In some sense, intending to E is itself taking means to E-ing. Having intended to attend the New York conference, I can realize the end of attending the New York conference with one fewer step than I can realize the end of attending the Boston conference. Boston requires, whereas New York does not, that I take the step of changing my mind.

This line of thought would support something stronger than IE1, in two respects.

Intention Effect, second pass (IE2): When there is *reason* for me to E, if I intend to E, then I have stronger reason, *on balance*, than I would otherwise have *to E* and to take some sufficient means to E.

First, intending E makes it the case that one has stronger reason, on balance, not only to take *means* to E, but also to *E itself*. The cost of E-ing as a whole is lower than it was, and so, to that extent, it has become an even better bargain. Second, the effect is not limited to underdetermined cases. Suppose the New York conference was the only option, so that I had conclusive reason to attend it. It would still be the case that, having decided to attend, I no longer need to take that step. To that extent, then, the prospective costs of attending are lower than they were, and so, again, it is an even better bargain that it was. Even if I lacked sufficient reason to attend New York, the lower prospective cost of attending would give me greater, although perhaps still not sufficient, reason, on balance, to attend New York.

This account is compatible with nonpsychologism. The relevant principle would be:

Cost: If some condition lowers the cost, going forward, of achieving some end, E, then, because of this, there is less reason *not* to E. As a result there is more reason, on balance, to E and hence, by FP, to take some sufficient means to E.

In cases in which intending to E is necessary for E-ing and carries some cost, intending to E lower the cost, going forward, of achieving E, and so gives me more reason, on balance, to E. So (i) is satisfied. (ii) is clearly satisfied. And (iii) is satisfied, since conditions other than my having some attitudes can lower the cost, going forward, of achieving some end. The thing that I intend to buy, for example, can be put on sale.

This account may be correct, but it depends on two uncertain claims. Observe that the cost of changing one's mind, in itself, is not very great. If one must deliberate further in order to change one's mind, then perhaps it might be more costly.¹¹ But, having already settled that Boston is as good as New York, one need not deliberate further. One can just change one's mind. The point is not that one is *likely* to change one's mind. Intentions are generally stable. And to the extent that one doesn't deliberate further about the decision, the option of changing one's mind is unlikely to suggest itself. The point is simply that, in fact, one has the option of changing one's mind, and taking this option is not particularly costly.¹² The first uncertain claim,

¹¹Scanlon observes that, because of the cost of deliberation, deciding to E, after giving due thought to the matter, usually provides one with "second-order purely pragmatic" reasons not to reopen the decision, unless some evidence appears that casts doubt on the choice that one made. ("Reasons: A Puzzling Duality?" 240. He credits it to Parfit, who in turn credits it to Broome.) If one has given the matter *due* deliberation, then there is no reason to expect that more deliberation will deliver a better result, and there is reason to expect that it will waste one's time and attention. As Scanlon notes, however, this means only that one lacks sufficient reason to *deliberate further* about which end to intend, not that one has greater reason to take steps to the end that one intends.

¹² In conversation, Michael Bratman suggests that this may not always be so. Suppose, as is often the case, that the option, E1, that one presently intends is better coordinated with one's other worthwhile intentions than the other option, E2, would be. In order to change one's mind to E2, one must revise one's other intentions that are coordinated with E1 and replace them with intentions that are coordinated with E2. This requires further deliberation, and so it carries a nonnegligible cost. While there is something right about this suggestion, the focus on the cost of deliberation is somewhat beside this point. The basic insight is simply that intending E1-coordinated actions gives one stronger reason to intend to E1 than to E2, because it makes one's intention to E1 more likely to succeed than one's intention to E2. (The cost of deliberation is relevant, if it is, only to a further question. One could change one's situation so that one had

then, is that such a small reduction in the cost of New York makes it the case that one lacks sufficient reason to take the means to Boston. In defense of this claim, one might urge that there is a perfect balance, so any reduction in cost tips the scales. But it is not clear that, in judging that the choice is underdetermined, we must assume so fine a balance.

Now suppose that intending New York does tip the scales, making it the case that one lacks sufficient reason to attend the Boston conference. Given that one has conclusive reason to attend one of the two conferences, it follows that one has conclusive reason to attend New York. Therefore, presumably, one has conclusive reason to sustain one's intention for New York. Therefore, if one changes one's mind, even with perfect dispatch, then one defies a conclusive reason. This is the second uncertain claim, which many appear to deny.¹³

3. If we likewise deny that, if one changes one's mind, even with perfect dispatch, then one defies a conclusive reason, then we accept, in effect, a constraint on IE1:

Intention Effect, third pass (IE3): If I have merely sufficient reason to E, then if I intend to E, I have stronger reason than I otherwise would to take *means* to E-ing. However, my intending to E has no comparable effect on my reason to *E itself*.

Our question is then how to explain IE3.

greater reason to intend to E2 than to E1: namely, by replacing one's E1-coordinated intentions with E2-coordinated intentions. Does one have reason to change one's situation in this way? Well, one reason against doing it is that it would require costly deliberation.) It is also worth noting that this suggestion does not help to explain IE. If one intends E1-coordinated actions, then intending to E1 will be more effective, so one has more reason to intend to E1. But this does not explain why intending to *E1*, as opposed to intending *E1-coordinated actions*, affects one's reasons: why it should itself give one stronger reason to E1 or take means to E1. All the same, the suggestion identifies a different mechanism by which one's intentions affect one's reasons, which I discuss in section 5, below.

¹³Broome, "Are Intentions Reasons?"; Raz, "The Myth of Instrumental Rationality," p. 22; Scanlon, "Reasons: A Puzzling Duality?"

We can begin by observing that means are often conditional. Quite often, the *effectiveness* of M-ing as a means to E—that is, the degree to which M-ing raises the probability that one E’s—varies with some further condition, C. At the limit, M may not raise the probability that one E’s at all unless C obtains. In this limiting case, if C does not obtain, then the fact that M-ing *would* help one to achieve E, if C *did* obtain, is not a reason to M. If I don’t have a hundred million dollars, for example, then it is not a reason to raise my paddle (in response to the auctioneer’s, “Do I hear one-hundred million?”) that raising it is necessary to acquiring the Picasso on auction. It is true that I won’t acquire the Picasso unless I raise my paddle, but this is not a reason to raise my paddle, because it is also true that I won’t acquire the Picasso even if I do raise it. Raising my paddle is not an *effective* means, or a means *strictly speaking*; it does not, as things are, raise the probability that I acquire the Picasso.

Observe, next, that often one such condition is that one takes *other* means. If M1-ing is *not sufficient* means to E-ing, then M1-ing raises the probability that one E’s only to the extent that one is likely to do something else: say, M2. Registering does not raise the probability that I attend the conference if there is no chance that I will fly to New York.

Observe, finally, that, quite often, one is likely to take other means, such as M2, if and only if one intends to E. If we attempted to explain this observation by claiming that intending to E gives an agent stronger reason to M2, then our account would be circular. We are trying to explain why one acquires, by intending something, reason to take means to it. But we do not need this normative claim to explain the observation, and it unclear how any normative claim could explain it. The observation is explained instead by nonnormative, causal regularities between an agent’s intending an end and his taking means to it.

Let us put these pieces together. Suppose that M1-ing is some means to E, but not sufficient means. M1-ing and M2-ing together are sufficient means. Even if one M1's, one will not E, unless one also M2's. One will M2 if and only if one intends to E. If one does not intend to E, then one will not M2. So, M1-ing will not make it more likely that one E's. So, it is not a reason to M1 that it makes it more likely that one E's. By contrast, if one intends to E, then one will M2. So, M1-ing will make it more likely that one E's. So, it is a reason to M1 that it makes it more likely that one E's.

Returning to our example: I decide to attend the New York conference and not the Boston one. The Boston organizers ask whether I wish to pre-register, indicating that it's cheaper than registering on site. Do I have sufficient reason to pre-register? No. So long as I have not adopted the goal of Boston, it's all but certain that I will not take other means that, along with registration, would be jointly sufficient to attending the Boston conference. I won't book a flight to Boston, I won't board a Boston-bound plane, etc. So pre-registering would be idle. Now the New York organizers call, to make a similar proposal. Do I have sufficient reason to pre-register for New York? Yes. Since I intend to attend the New York conference, it is probable that I will take other means that, along with registration, will be jointly sufficient to attending the New York conference. I will book a flight to New York, etc. So pre-registering would not be idle. It would contribute to making it the case that I take *sufficient* means to New York, which I have sufficient, but not conclusive, reason to do—sufficient reason that was there before my decision and that my decision did not change. So I now have, as a result of my

decision, sufficient reason to pre-register for New York, whereas I still lack sufficient reason to pre-register for Boston.¹⁴

On this account, intending to E does not affect one's reason to E, and so does not affect one's reason to *sustain* the intention to E, or to take *sufficient* means to E: e.g., to do *all* of: M1-ing and M2-ing. It affects only one's reason to take *insufficient* means to E: e.g., to M1 on its own (or, *mutatis mutandis*, to M2 on its own). So what we have is not quite IE3, but instead:

Intention Effect, fourth pass (IE4): When there is merely sufficient reason for me to E, if I intend to E, then I have stronger reason than I would otherwise have to take *insufficient* means to E-ing. However, my intending to E has no comparable effect on my reason to E.

Does this leave a fact to be explained: namely, that intending the end has a comparable effect on my reason to take *sufficient* means?¹⁵ It is difficult to see how there could be such a fact to explain—how intending the end could have a comparable effect on one's reason to take *sufficient* means to E—if, as the constraint says, intending the end has *no* comparable effect on my reason to E. Why? For one thing, E-ing is sufficient means to E-ing. For another, suppose that there is no sufficient means to E that has any “side benefits.” The only reason to take some sufficient means is to bring about E, and presumably, one has as much reason to take some sufficient means as one has to E. If intending to E gave one greater reason to take some *sufficient* means to E-ing, but intending to E did *not* give one greater reason to E, then we would get the bizarre

¹⁴ In the context of a broader exploration of the normative significance of intending, Brunero, “Are Intentions Reasons?” arrives independently at a similar explanation. I have learned a great deal from his paper.

¹⁵ In conversations about this paper, I have been asked: “What about ends that are achieved by *noncomposite* means: means that are not made up of separate parts?” Since such noncomposite means would have to be *sufficient* means, this possibility is addressed in the present discussion.

result that one had greater reason to take some sufficient means to E-ing than to E itself, even though the only case for taking those means is to bring about E.

Although I have been focusing on underdetermined decisions, decisions of other kinds can have this effect on reasons to take means. Suppose that I have conclusive reason to attend the New York conference, but I haven't decided to. Do I have sufficient reason to pre-register? Not obviously. If I don't decide to attend the New York conference, then I won't take other means. So I won't attend the New York conference even if I pre-register. So whatever reasons I might have for attending it do not translate into reasons for pre-registering. And there may well be other reasons against pre-registering, such as that the money could be spent elsewhere. Thus, while the best option may be to pre-register and to take other means, the second-best option, assuming that I will not take other means, is not to pre-register at all.¹⁶ Furthermore, even though I lack sufficient reason for Boston, if I am intent on it, I may have more reason to take insufficient means to Boston than to take insufficient means to New York. Again, it may be the second-best option.

What we have, finally, is:

¹⁶ This is contrary to the suggestions of Parfit, Raz, and Scanlon that it is only in underdetermined cases that intending an end affects one's reason to take the means to it. Nevertheless, there may be something to the idea, even on this account, that this effect is more pronounced in underdetermined cases than in determined cases. In most cases, one is more likely to decide to do something that one hasn't yet decided to do if one has conclusive reason to do it than if one has merely sufficient reason to it, other things equal. Deciding to do something that one has conclusive reason to do therefore has a smaller effect on the probability of one's taking other means than deciding to do something that one has merely sufficient reason to do. And in most cases, one is less likely to carry out a decision to do something that one lacks sufficient reason to do than to carry out a decision to do something that one has sufficient reason to do, other things equal. Deciding to do something that one lacks sufficient reason to do therefore has a smaller effect on the probability of one's taking other means than deciding to do something that one has sufficient reason to do. In both cases, therefore, intending an end does less to strengthen one's reason to take insufficient means in determined cases than in underdetermined cases.

Intention Effect, final pass: If there is reason for me to E, then, if I intend to E, then I have stronger reason than I would otherwise have to take insufficient means to E-ing.

However, my intending to E has no comparable effect on my reason to E.

How is this compatible with nonpsychologism? What we need, recall, is some principle of the form:

If some proposition, p, of some type P is true, then, because of this, there is more reason for me to respond in some way R(p), where R is a function from propositions to responses,

where:

- (i) that I intend to E is, or makes true, some p of P
- (ii) $R(p) = \text{to } M$
- (iii) there is some proposition, p*, of P such that for no attitude A, p* is that I have A.

We have been implicitly appealing to such a principle: namely,

Effectiveness: If some condition makes my taking some step, M, a more effective means to some goal, E, that there is reason for me to pursue, then, because of this, there is more reason for me to M.

We have seen how (i) is often satisfied. Often, intending to E makes it more likely that one will take other means, and the more likely it is that one takes these other means, the more effective M is as a means to E. (ii) is obviously satisfied. And (iii), the crucial condition for nonpsychologism, is also satisfied. All kinds of conditions that do not consist in my having some attitude can make M-ing a more effective means to some end that there is reason to pursue.

4. Now I turn to some objections. On this account, the crucial thing that intending the end does is to make taking some insufficient means more effective: that is, to increase the degree to which taking some insufficient means raises the probability that one achieves the end.¹⁷ This probability would appear to be epistemic. Some might protest that while epistemic probabilities may affect one's reasons for *belief*, they cannot affect one's reasons for *action*. Suppose that intending to E has the described effects on the epistemic probabilities, but the improbable thing occurs: one does not take other means with which M1-ing would be jointly sufficient. Looking back, we should say that one did not have reason *to act: to M1*. Given how things would turn out, it was waste.

I am inclined to think that one did have reason to M1: that reasons for action depend on what is likely to happen, rather than on what in fact happens. But I do not need to press this point. Even those who insist that one did not have reason, strictly speaking, to M1 still think that, in some other sense, one's M1-ing was appropriate. It was, say, "reasonable" to M1, even if it turns out that one lacked reason, strictly speaking, to M1.¹⁸ If one chooses to speak in this way, then one should interpret my claim as that intending an end makes it more reasonable to take insufficient means to it, not that it gives one greater reason to take those means. Although I am inclined to think that a stronger conclusion is warranted, I do not think that it matters much for present purposes.

¹⁷ This is very similar to what Scanlon, in another context, calls the "predictive" significance of intention. See his "Moral Assessment and the Agent's Point of View," unpublished.

¹⁸ While it may be tempting to identify what is most "reasonable" (relative to the evidence) for one to do with what one has most reason *to believe* that one has most "reason" (not relative to the evidence) to do, this will not work. Derek Parfit's example of the trapped miners in *Reasons and Persons* (XXXX) represents a case in which one has incontrovertible reason to believe that one does *not* have most "reason" to do what would be most "reasonable" to do.

Of course, it need not always be true that intending the end makes insufficient means more effective. On the one hand, one's intending the end might not make it significantly more likely that one takes other means, because, say, one will refuse to take those other means. On the other hand, intending the end might not make it significantly more likely that one takes other means, because it was already so clear that one would take those means. In these cases, according to the proposal we have been considering, one's intending the end does not affect one's reason to take the means in the way it otherwise would.

Viewed from the outside, this seems true and unproblematic. If we observers believe that the agent won't ever board the plane, say, then it will seem to us a waste for him to pre-register. And if we are convinced that he will opt in the end to go to the New York conference, however sincere his professions of indecision, then it will seem to us that he ought to pre-register now, before the fees increase. It will probably seem to us that he ought to decide to attend now too; he's wasting his time turning it over in his mind, simply forestalling the inevitable.

When we consider the agent's point of view, however, things may be more puzzling. It would be strange for such predictable agents to view their own intentions for the end as making no difference to whether they should take the means. On the contrary, we would expect each, upon making his decision, to see himself as having reasons that he did not have before. The explanation, I think, is that it would be difficult or impossible for the agent to believe what the observers know, so long as his intentions are as described. If the first agent knew what observers know—namely, that he will not take other, necessary means—then he could not sustain his intention for the end. One does not intend what one believes one will not do.¹⁹ If the second agent knew what observers know—namely, that he will decide on the end—then he would likely

¹⁹See Bratman, *Intention, Plans, and Practical Reason*.

feel pressed to decide on it himself. Once one knows how one will decide, it is bound to strike one as idle to continue wrestling with it as though it were a live question.

This broaches a more general, if somewhat inchoate, worry about this account. The account may seem to suggest that agents must take oddly “theoretical” or “predictive” stances toward their own intentions and actions: that agents must treat what they will or do as though it were a fact about their situation with which they were simply confronted. This is a worry that applies even in what ought to be the standard cases for the account: cases in which intending the end actually does make it more likely that one will take other means to it.

It is worth noting that this worry arises only when we consider these judgments about reasons from the first-person, deliberative point of view. On the one hand, this makes the worry worth taking seriously, because reasons are, so to speak, made for the deliberative point of view. On the other hand, claims about reasons can nevertheless be made from other points of view. And when observers advise or report, in the second- or third-person, that, because of what an agent does or does not intend, he has stronger or weaker reason to do or intend something else, such advice or reporting, as noted earlier, seems true and unproblematic. This encourages me to think that the worry about an oddly “theoretical” stance is, in the end, superficial.

Suppose I reasoned: “Because I have decided to go to the New York conference, I will almost certainly be there. In that case, pre-registering would not be a waste. So it makes good sense to pre-register.” What about this line of thought would be oddly “theoretical”? It might be said: “In order to reason from the content of my expectation about what I will do to a belief about what else I have reason to do, I need to take *what I will do as given*, just as I might take any other circumstance as given. It would be oddly ‘theoretical’ to take what I will do as given, and not within my control.” This line of thought, however, equivocates on “take as given.” I

must take it as given that because I have decided to do it, I will (if all goes well) do it. I need not believe that what I will do is somehow beyond my control. I can know that I had the power to decide otherwise than I did, while also knowing that I decided as I did, and I can know that I still have the power to change my mind, while remaining self-consciously settled in my present course.

It might be suggested that *any* reasoning from the content of my expectation about what I will do to a belief about what else I have reason to do would be oddly “theoretical.” But, first, such reasoning is pervasive. Should I take the mail on this outing? Well, it depends on where I will be going. If I will be passing the mailbox, then yes; if I will not, then no. Second, this implicit ban on reasoning from expectations about what we will do would sequester our practical deliberation in an unnatural way. We would have expectations about what we will do and what will happen when we do it, on the one hand, and expectations about what the future will bring independent of what we do, on the other. While we would be able to deliberate theoretically from expectations of both kinds, we would be able to deliberate practically only from expectations of the latter kind. Finally, one might wonder how we could ever take means to our ends, or coordinate our actions, if we could never deliberate from expectations about what we will do. It might be answered that we accomplish this by directly conforming to rational requirements on our intentions, such as Means-End, in a way that bypasses expectations about what we will do. When we intend the end, we come to intend the apparent necessary means, without any reflection on what we are likely to do. But this does not seem enough. For one thing, beliefs about how likely I am to succeed are relevant to my deliberations. The less likely I am to succeed, the less reason I have to invest in the means. If we reason only from the content

of our intentions (whatever, precisely, this would be), then we cannot be sensitive to such reasons.

Perhaps what seems oddly “theoretical,” then, is the suggestion that I reason from the content of my awareness that I *intend* something to an expectation that I will *do* certain things because of it. My awareness of what I intend usually makes an implicit and immediate contribution to my expectation of what I will do, even if reasoning from the contents of beliefs about other relevant circumstances may strengthen or weaken that expectation. Thus, it would be more natural to begin with the thought: “Because I’m *going* to go to the New York conference...” than to begin with the prior thought: “Because I *am in the state of intending* to go to the New York conference...” All of this seems right. The question is why it threatens the proposal. It would threaten it if we accepted some thesis to the effect that a fact affects our reason for some attitude only if, when we adjust our attitudes accordingly, we normally explicitly reflect on its having that effect. Yet this would be implausible as a general thesis. My having a perceptual experience as of something’s being so affects my reason to believe that it is so. But in the vast majority of cases, I do not explicitly reflect on my perceptual experience itself, or on its having an effect of this kind.

5. I have been trying to show that if we are nonpsychologists, who hold that our attitudes do not, as attitudes, provide us with reasons, we can still account for the familiar ways in which having attitudes, or at least intending something, changes our reasons. What I want to mention in closing is the relevance of this conclusion to principles of rationality, such as Means-End.

Why ought we to conform to principles of rationality? Whence their “normative force”? One proposed answer is instrumental. Conforming to rational requirements as a rule leads us,

over the long run, to attitudes that accord better with reason.²⁰ If our means-end beliefs are reliable, for example, then conforming to ME leads us to more beliefs and intentions that we have conclusive reason to have, and fewer beliefs and intentions that we lack sufficient reason to have. The answer that I favor, by contrast, is that the normative force of the principles of rationality is a persistent and abiding illusion. This “Transparency Account,” as I have called it elsewhere, rests on the claim that putative principles of rationality such as ME are approximations to, or special cases of, a more fundamental principle: to follow one’s judgment about one’s reasons. If what rationality requires of one is that one follow one’s judgment about one’s reasons, then it will always *seem* to one, when rationality requires a response of one, that one has reason to give that response—for that is precisely what one has already judged.²¹

I will not try to decide between these views here. My present point is that both views, when applied to ME, would seem less plausible if the following claim were false:

Means-End Benefit (MEB): If I violate ME and my means-end belief is true, then there is some change that ought to make. That is, if I intend at t to E, believe, truly, at t that I E only if I intend at t to M, and do not intend at t to M, then either I lack sufficient reason to intend at t to E, or I lack sufficient reason to believe at t that intending at t to M is a necessary means to E-ing, or I have conclusive reason to intend at t to M.

If MEB is false, then why should conforming to ME be expected to lead one to fewer attitudes than one lacks sufficient reason to have and to more attitudes that one has conclusive reason to have? And why should one judge, of some way of conforming to ME, that one ought to conform in that way? The difficulty is that, unless intentions affect our reasons, MEB is not true. There

²⁰ See Bratman and Broome. I criticize this suggestion in “Why Be Disposed to Be Coherent?” (draft).

²¹ “Why Be Rational?” *Mind* XXXX.

are four relevant possibilities (dropping the temporal markers): (a) I lack sufficient reason to believe that intending to M is a necessary means to E-ing; (b) I lack sufficient reason to intend to E, (c) I have conclusive reason to intend to E, and (d) I have merely sufficient reason to intend to E. If I lack sufficient reason (a) to hold the means-end belief, or (b) to intend to E, then the consequent of B is true. I have (c) conclusive reason to intend to E only if (ignoring “state-given” reasons) I have conclusive reason to E. If I have conclusive reason to E, then (plausibly) I have conclusive reason to take the necessary means to E-ing, which includes intending to M. Therefore, if I have conclusive reason to intend to E, the consequent is true. So far, so good. But what if (d) I have merely sufficient reason to intend to E? Why should it follow that I lack sufficient reason to hold the means-end belief, or that I have conclusive reason to intend to M? I might well have sufficient reason to hold the means-end belief and merely sufficient reason to intend to M. If so, then I can intend to E, but not intend to M, and yet there might be no change that I ought to make.

We avoid this conclusion if we grant that intentions do change our reasons in the way that we have been considering. My only reason to intend to E is that it makes it more likely that I E. If intending to E does not make it more likely that I E, then I lack sufficient reason to intend to E. And if I do not intend to M, then intending to E does not make it more likely that I E. Therefore, insofar as do not intend to M, I lack sufficient reason to intend to E. Just as intending the end can give one greater reason take the means, failing to take the means can give one less reason to intend the end. This is another application of Effectiveness. Intending to E is itself insufficient means to E-ing. I must also take other means: in particular, intend to M. To the extent that I do not these take other means intending to E is not an effective means. This point guarantees that the consequent is *always* satisfied, *whatever the reason* to intend to E or to intend to M.

However, while I lack sufficient reason to intend to E, given that I do not intend to M, I may have sufficient, or even conclusive reason to intend to M, given that I do intend to E. Indeed, this may be the first best. If this is the case, and if I intend to M, then I may no longer lack sufficient reason to intend to E. This is another way out of my predicament.

In this paper, I have been discussing mainly normative relations between ends and *means*: that is, *doing* what makes a given end *more* likely to be achieved. But there are also normative relations having to do with *coordination* among ends: that is, *refraining* from pursuing other ends that make a given end *less* likely to be achieved. Suppose that ends E1 and E2 are otherwise comparable, except that E1 is better coordinated with F1, whereas E2 is better coordinated with F2. This is just to say that if one intends to F1 but does not intend to F2, then intending to E1 will be more effective than intending to E2, and vice-versa. By Effectiveness, one has greater reason to intend E1 than to intend E2. This represents another way in which one's intentions can affect one's reasons. In this case, what happens is not that intending an end affects one's reason to *achieve that end*, or to *take the means* to it. Instead, intending one end affects one's reason to *intend something else*.

This observation is relevant to attempts to explain the normative force of another putative principle of rationality:

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

As before, both the instrumental justification and the Transparency Account would seem less plausible if the following principle were false:

Consistency Benefit (CB): If one believes, truly, at t that if one X's then one does not Y, intends at t to X, and intends at t to Y, then there is some change that one ought to make:

i.e., either one lacks sufficient reason to believe at t that if one X 's then one does not Y , or one lacks sufficient reason to intend at t to X , or one lacks sufficient reason to intend at t to Y .

As before, unless intentions change our reasons, CB is false. Suppose that I have sufficient reason to believe that if I X then I do not Y and sufficient reason to intend to X . Why should we expect that I lack sufficient reason to intend to Y ? Suppose that I know that I cannot attend both New York and Boston. Can't it still be the case that I have sufficient reason to intend New York and sufficient reason to intend Boston? Isn't that the very situation that we have been considering? Why then should CB seem so much as plausible?

We need to take into account my other intentions. We must consider whether, *if I intend to Y*, I have sufficient reason to intend to X . When I cannot both X and Y , often (although not necessarily²²) my intending to X is uncoordinated with my intending to Y , and vice-versa. As a result, if I intend to Y , then my intention to X will be less likely to succeed. By Effectiveness, therefore, if I intend to Y , then I have less reason to intend to X . This helps to explain how, on the one hand, I can have sufficient reason to intend New York if I do not also intend Boston, but, on the other, lack sufficient reason to intend New York if I also intend Boston. While this shows why CB may seem plausible, it does not guarantee that CB always holds. This is, I think, as it should be. CB, and indeed Consistency itself, are often, but not always, true. But this is a topic for another paper.

²² I discuss exceptions in "Why Have Consistent Intentions?" (draft).