

## Consequents of True Practical Conditionals Detach<sup>1</sup>

(4388 words exclusive of footnotes and bibliography, give or take a word or five.)

When people want to capture the demands of instrumental reason or other hypothetical norms it is natural to express these claims in the form of a conditional of the form, ‘If you want to  $\Phi$ , then you ought to  $\Psi$ ,’ or ‘If you believe you ought to  $\Phi$ , you should  $\Phi$ . There are of course debates about which hypothetical norms in fact apply independent of any worries about such conditionals, but recently at least some debates have been occasioned by worries about the conditionals themselves. Some seemingly true instances of such conditionals with true antecedents have unpalatable consequents. ‘If someone wants/intends to kill his father in an undetectable way, he ought to poison his father’s food,’ could be recommended by instrumental rationality if poison is undetectable. When applied to patricidal but prison wary Harry, the conditional seems to entail that Harry ought poison his father. Yet Harry ought not poison his father. Similar examples proliferate using various more abstract principles of instrumental rationality connecting perceived efficient means to desired ends in general. In each case we are hesitant to detach the consequent of the conditional when that consequent is abhorrent or otherwise questionable. Many philosophers now think that such arguments must thus be invalid.

Two strategies have presented themselves to avoid detaching in such situations. One is to deny the acceptability of conditional oughts with antecedents that don’t have positive normative status. (Korsgaard 1997; Hampton 1998) More popular yet has been to take ‘ought’ as a sentence operator which takes wide-scope over the whole conditional (Hill 1973, Greenspan 1975,

---

<sup>1</sup>So far already conversations with Hagit Benbaji, David Enoch, John Gibbons, David Heyd, Nadeem Hussein, Matt King, Jennifer McKittrick, Simon Rippon, Mark Schroeder, Sigrun Svaversdottir, Ralph Wedgwood and Ruth Weintraub have been helpful in writing this.

Broome 1999, 2001, 2002, Wallace 2001). Applied to the above example, the suggestion is that instrumental rationality requires that Harry not both want/intend to undetectably kill his father and not use poison to do so.. And this requirement can be met by not wanting/intending to kill his father. The consequent does not detach to yield a free-standing ought claim because the conditional is true so long as Harry ought not want/intend to kill his father.<sup>2</sup>

The wide-scoping strategy is reinforced by the well-earned prestige of its main proponents and by the support it gets from treating ‘ought’ analogously to terms like the necessity operator of standard model logic. Yet, I think it is wrong. And I think it is wrong for a reason that also counts against the first avoidance strategy. When we have true conditionals of instrumental rationality even with abhorrent consequents and when the antecedents are true, the consequents are also in fact true. In other words, the consequences do in fact detach. Our understandable hesitance to assent to the detached consequents does not indicate that they are false when correctly interpreted. Rather, I argue, the hesitancy stems from a tendency to take free-standing assertions of these consequents as expressing something they do not.<sup>3</sup>

My main argument has two parts. The first proposes an analysis of ‘ought’ claims as reason claims and suggests that pragmatic considerations on the use of these sentences are capable of explaining why it is often infelicitous to assert their consequents detached from the

---

<sup>2</sup>Stephen Finlay (forthcoming) was helpful in writing the setup for the paper here.

<sup>3</sup>Conditionals with normative consequents are abundant and not all of them are true. I don’t have a general story about which ones are. I do think there is at least one norm of instrumental reasoning sufficient to generate some true practical conditionals of the sort that has been the prime focus of this debate. But I don’t know precisely what the appropriate principle of instrumental reasoning is, whether it governs desires in general, only serious desires, intentions or whatever. I’m hoping I can discuss the general issue of detachment without picking a favorite.

antecedents, even when they are true. Thus I offer the resulting package as a candidate for the best explanation of the intuitive plausibility of the conditionals together with the intuitive inappropriateness of free-standing assertions of their consequents. But this package requires that there be reasons of the right sort for the claims in question to be true when I say they are. Thus the second component of my argument is a defense of the claim that there are such reasons.

### **Component One: Ought claims as Reason Claims**

I'm proposing that the contents of 'ought' claims should be cashed out as claims about reasons. The claim that Harry ought to  $\Phi$  has the content that Harry has a (pretty good) reason to  $\Phi$ . For what it is worth, I think that this claim says that Harry stands in a certain relation to an act type,  $\Phi$ -ing.<sup>4</sup> This relation is grounded in certain features of Harry's situation, and at least some of these features count as the grounds of his reason. Sometimes those grounds include facts about Harry and his psychology and they may include much more besides these. However, the claim that Harry has a reason to  $\Phi$  does not by itself tell us what those grounds are, though we often manage to communicate those grounds when we talk about Harry and his reasons. We could say that Harry has a reason to  $\Phi$  because  $\Phi$ -ing would make him healthy, or because it would help his friend Lisa, or whatever.

Reasons can be prima facie or they can be all things considered. Very often our situation is such that we have different conflicting reasons to do incompatible things. Usually the conflicting reasons are grounded in different facts about the agent's situation and different reasons to do the same thing might depend on different sets of grounds as well. The resulting

---

<sup>4</sup>Harman (1973) Geach (1981).

reasons interact (in possibly very complex ways) to determine what we have all-things-considered reason to do. To say that Harry has reason to  $\Phi$  does not yet tell us whether the reason is prima facie or all things considered. Philosophers of course have coined the terms ‘prima facie’ and ‘all-things-considered’ precisely to make explicit which are which. And we already had ways to express the distinction using somewhat awkward but perfectly good locutions such as, ‘Insofar as Harry wants to go to the store, he ought to . . . ,’ and ‘Given the situation as a whole . . . ,’ and so on. But the bare claim that someone has reason to  $\Phi$  does not tell us whether she has all-things-considered reason or only prima facie reason. (Note that when someone has an all-things-considered reason to  $\Phi$  they will always also have prima facie reason to  $\Phi$ .)

Typically we convey the grounds of an agent’s reason or possible reasons when we embed a reason claim in a conditional. ‘If Mary wants to stay dry, she has reason to take her umbrella,’ surely gets across quite a lot about Mary’s reason to take her umbrella. My proposal is that ought-claims work the same way. To say that Mary ought to  $\Phi$  says just what the corresponding reason claim says. Similarly ‘Mary ought to take her umbrella if she wants to stay dry’ conveys just the same information about what grounds those reasons, as the similar conditional claim that she has reason to take the umbrella. This will be important later.

If this is the right story to tell it will turn out that the consequents of true conditionals of this sort will detach at least where we quantify over all reasons and not just over all things considered reasons. If it really is true of Mary that if he wants to stay dry she ought to take her umbrella, and it is also true that she wants to stay dry then it will also be true that she ought (at least prima facie) to take an umbrella. By parity of reason, the proposed account says the same

thing about Harry and his homicidal goals. If the conditional relating Harry's hypothetical desire to kill his father to using poison is correct, then should he in fact want to kill his father he will have a prima facie reason to use poison. So a free-standing instance of the consequent is true.

#### Possible Subhead Here

What of the intuitively plausible claim that he ought not poison his father? This is likely true as well. Most children have very good reason not to poison their parents and we all have good reasons not to kill another human being. Thus Harry likely has stronger reason not to kill his father and also stronger reason to rid himself of his homicidal desires, but that is consistent with the positive claim that as things stand he has a reason to use poison. On the hypothesis that oughts judgements express reason claims, and that many of these claims are prima facie, there is no problem with the thought that one both ought and ought not do one and the same action or type of action. It is the essence of prima facie reasons that they can conflict and recommend incompatible courses of action. So there is no logical trouble here so far.

But many people have the intuition that the free-standing occurrences of claims such as the consequent of the conditional about Harry are *false*, and *hence* inappropriate. My suggestion is that they think that these claims are false precisely because they sense that they are *not assertable* but that their inappropriateness is due *not to falsity* but to the misleading nature of the bare assertion where the proximity of the antecedent has not made salient the prima facie character of the reason that makes it true.<sup>5</sup>

One natural home for ought claims and practical conditionals embedding them is in

---

<sup>5</sup>Most of what I say here is an elaboration of the basic idea from Schroeder (2005, p. 8; 2008a, p. ?) that pragmatic implications can explain much of our resistance to assert the consequents of such conditionals.

asking for and giving advice. This generates a background of conversational purposes which allow Gricean mechanisms for generating conversational implicatures to affect our judgements of what it is appropriate and inappropriate to say.<sup>6</sup> As is well known, Grice suggested that much of what we communicate when we talk is conveyed through conversational implication as opposed to by the literal content of what is said.<sup>7</sup> He suggested that this was possible because we normally assume that our conversational partners are cooperating with us to communicate. Cooperative speakers will try to hold to maxims requiring conversational contributions to be as informative as needed for purposes of the conversation but no more so, and also to make contributions relevant to these purposes. And because they do this we can infer information beyond the semantic content of an utterance from the fact that it and not something else was what was said by the speaker.

If we are giving advice about what to do or think, more important reasons are more relevant to the purposes of the conversation than less important reasons. All things considered reasons are, by their nature the final determinants of how it is rational to act and hence of the highest relevance in determining what to do. Prima facie reasons are more important to the extent that they determine what all things considered reasons one has to act, and less important to the extent that they should have no impact or no significant impact on the advisees all things considered reasons.

According to the Gricean picture, when we are told that Harry ought to do such and such

---

<sup>6</sup> In defending similar ideas Mark Schroeder points out that intuitions that someone has no reason of a certain sort are unreliable because talk of reasons generally implicates that the reasons in question are pretty good. (Schroeder, 2005, p. 8; 2008a, p. ?)

<sup>7</sup>Grice, (1989) p. 28.

with no further information given to us, we are entitled to assume both that this is as informative as it need be for our purposes, but no more so and that it is relevant to our conversational goals. If this reason mentioned were but a very weak prima facie reason it would not be much help in advising Harry. If it were a stronger prima facie reason it would be helpful, but not as helpful as it would be with information about the other relatively strong prima facie reasons to be weighed against it. Furthermore, it would be helpful to know what grounded the reason in question rather than just that there is some relatively strong prima facie reason. And in normal cases if one knows there is a relatively good prima facie reason for Harry to act, one will also know its grounds. In these conditions a reasonable audience would infer that the information is as helpful as the speaker is capable of making it, and that the reason in question was an all things considered reason. For if it was a weak prima facie reason it would not further our conversational purposes. And if it was a strong prima facie reason the speaker would likely know more and tell the audience what that more was. In general, the point of telling someone that they have a prima facie reason will be best served if we also convey the grounds of the reason. Thus assuming the free-standing assertion of the conditional to be made in a cooperative spirit entitles the audience to assume the reason under discussion to be an all things considered reason, absent indications that it is not. But interpreted as all-things-considered claims, both the reason claim recommending the use of poison and the corresponding ought claim will be false. So there should be no surprise that free standing ought and reason claims of this sort sound false to competent listeners. The most natural interpretation of what they are meant to convey is false.

Of course this holds only of free-standing assertions of detached consequents when the

surrounding conversation does not fill in the requisite information to allow them to be interpreted as prima facie reason claims. Conversational implicatures are cancelable by further things a speaker might say. Here is one way this can work. Normally conversational implicatures are generated when one openly violates one or more of the conversational maxims. The audience then has to figure out how the speaker can do that while still holding to a principle of cooperation. When a speaker says more immediately after saying something that would by itself violate a conversational maxim, the totality of what is said may no longer violate the maxim. An utterance which would otherwise be irrelevant or uninformative enough to generate an implication will no longer be so if the speaker goes on to fill in the rest of what the audience might want to know.

To my ear the implications at work to explain judgements that free-standing assertions of the consequents of practical conditionals are false are cancelable, as the theory would predict.<sup>8</sup> “Harry ought to poison the wine, but only insofar as he wants to kill his father undetectably,” removes the implication that the ought in question is all things considered.

So far I’ve focused on the giving of advice. What about the other natural home of reason and ought talk, in the domain of action explanation? In such contexts we care most about the reasons an agent acts on, whether they are actually what she had most all things considered reason to do or not. Thus an audience who hears a reason claim in this context will not be as entitled to assume that the reason in question is an all things considered reason even in free-standing detached uses. I think it fair to say that we don’t have nearly as much trouble accepting

---

<sup>8</sup>Again, Schroeder makes essentially the same point about reason claims. (Schroeder, 2005, 8; 2008, ?)



such free-standing consequents as parts of a rationalizing explanation as we do where we are worried about giving advice, though the issue is tricky.<sup>9</sup> If we are trying to explain why Harry poisoned his father, there are different ways to frame the issue. We might want an explanation of the whole course of action, we might want to know why he used poison (rather than a gun) or even why it was his father (as opposed to someone else). I'm not sure which of these is the right test case. Maybe they all are. Whichever we are interested in we might think that the claim that he had reason to do so is minimally informative as an answer. But to the extent this is a problem for the explanation it doesn't seem that the problem is one of falsity. And we can say more to make the answer better, for example by giving the reason. We could felicitously say "Harry had a reason to poison the wine insofar as he wants to kill his father undetectably." Or we might say, "He did what he ought to have, given just his goal of poisoning his father."<sup>10</sup> Or we might say, "He had a reason," and leave it at that. In this context, the fact that we said no more would not implicate that the reason was all things considered. It would more likely be taken to indicate that there was some reason not to say more. I'm a little confused about the fact that, "He did it

---

<sup>9</sup> I should admit that ought claims do not get used as naturally as reason claims when it comes to explanation. Since those who dislike detachment for oughts generally have the same attitude when it comes to detaching the consequents of claims about reasons, the asymmetry about reasons is at least some reason to worry about the strategy adopted by the enemies of detachment. Still, I should admit that the two locutions aren't entirely on a par when explanation is at issue.

<sup>10</sup> You might not think this is true. But that might be because you accept a particular conception of reasons which rules this out. It should be no surprise that people's judgements about what counts as cooperating or as violating a maxim depends in part on their background beliefs. If you think people's bad goals never give them reasons, the additional words that I claim cancel the implicature in these cases won't seem to you like they do the work precisely because you think what they say is necessarily false. This is why I think the second part of this paper, where I argue that we should accept that there are reasons of the right sort has to be part of a full defense of the view.

because he ought to have,” doesn’t sound to me like it works the same way. I hear this as an all things considered ought claim, but I can’t say why I should.

So I tentatively conclude that the correct semantics for ‘ought’ or ‘reason’ as used in the practical conditionals that motivate wide-scopers is one on which the semantic content of the consequent is true just in case the agent stands in the reason relation to the act-type mentioned, and that this relation can be either all things considered or prima facie. The reason we are often inclined to dissent from detached utterances of these true consequents is that their utterance often pragmatically conveys that the reasons in question are all things considered or very weighty when neither is true.

### **We Have Independent Reason to Think That Reasons of This Sort Exist**

This mixture of semantics and pragmatics as an explanation of the standard intuitions demands something of the metaphysics of reasons. It demands that the reasons postulated in fact exist. It would be nice to have some grounds independent of the very phenomena I’m trying to explain for thinking there are such.

Reasons are rationalizers, potential and actual. Roughly they’re facts which, if the agent were aware of them would rationalize whatever they are reasons for – an action if they’re practical reasons, a belief if they’re epistemic reasons. Sticking with the practical, different facts accessible to the agent might rationalize different actions, and agents might have reasons at various levels of accessibility to them which we might distinguish as more or less subjective and

objective reasons.<sup>11</sup> Whether they do so or not depends on still further facts that form the background conditions for their acting, as enablers, as the absence of defeaters or whatever.<sup>12</sup>

As Bernard Williams' observes, as rationalizers reasons must be able both to explain action and to justify or make sense of it.<sup>13</sup> This observation makes some sort of internalist constraint on reasons- one that ensures that reasons be able to motivate - highly plausible. As Robert Johnson puts it, "any account of what reasons are must make plain how, roughly, *the* reason there is for A to  $\phi$  could be A's *reason* for  $\phi$ ing. And something could be A's *reason* for  $\phi$ ing only if it could feature in both (i) the intentional explanation of his  $\phi$ ing, and (ii) his rational justification for  $\phi$ ing." (p. 59) The problem is to find the right way to capture this and to do so without causing problems for our judgements that real, less than ideal agents (agents like Harry) have reasons to do what they have reason to do.

A natural thought is that we have a reason to  $\Phi$  just in case we would  $\Phi$ , and  $\Phi$  on the basis of the reason in question, if we were informed of the grounds of the reason and completely rational. But this won't do for several reasons. The grounds of the reason may themselves be incompatible with full rationality. My psychological disorder may be a reason to see a psychologist, a reason I would not have were I fully rational.<sup>14</sup> And, insofar as some of my reasons are mere prima facie reasons, even when fully rational they might not motivate me to act

---

<sup>11</sup>van Roojen (2000) I want to remain neutral about whether the reasons are identical to the facts which rationalize the action or whether they are constituted by them. So my talk may seem sloppy at points since I have a tendency to slide around between the views.

<sup>12</sup>Schroeder (2008a)

<sup>13</sup>Williams (1981)

<sup>14</sup>Shope (1978), Johnson (1999), van Roojen (2000) and Sobel (2001)

because I have stronger prima facie reasons not to  $\Phi$ . Is there a way to capture the spirit of the natural thought motivated by the justificatory and explanatory role that reasons must be able to play?

The thought that makes rational idealization attractive is the thought that reasons must justify actions. If a rationally ideal agent would act on a reason this would seem to ensure that it makes sense to act on it. But we are blocked from taking that route by the examples above. We can make headway by recognizing that rationality comes in degrees. A consideration counts as a reason to  $\Phi$  if it would be more rational than not to act on it in the actual circumstances where it applies and other things are equal. And an agent has an all things considered reason to  $\Phi$  just in case it would be more rational than not to act on it in the actual circumstances of the agent.<sup>15</sup>

This isn't empty despite the remove from full rationality. We very often have a pretty good idea how a person who is in some way irrational might, consistent with that irrationality, behave more or less rationally. And we often have a pretty good idea of the grounds for these judgements in features of the person's situation and the rational requirements that interact with them. So, the suggestion is, in thinking about reasons we should be sensitive to what makes a person more or less rational, not just what they would do if they were completely rational.

Those of us who are unhappy with a wide-scope ought operator as a way of capturing the force of instrumental reasoning are unhappy with that suggestion precisely because, in the normal cases, even an agent's bad ends generate rational pressure to act in service of those ends. If the wide-scope strategy is successful at undermining detachment it won't adequately capture

---

<sup>15</sup>The need for the other things equal clause is to get prima facie reasons in, but it is unneeded if we are talking about all things considered reasons.

this pressure. The whole point of the wide-scoping strategy is to avoid the pressure to detach when the antecedent is fulfilled. Those of us who like detaching like it because we think reasons explain why, once the antecedent is fulfilled, it would be more rational than not to go on and do what the consequent suggests.

Barry intends to build a door. In order to build a door it would be very helpful to use a block plane and they are ready to hand. Barry knows this. This seems to underwrite the conditional claim that, "If Barry intends to build a door he ought to use a block plane." I'm claiming that the fact this represents is that if Barry had that intention he has a reason to use a block plane. And I'm claiming that when one has a reason to act there has to be some circumstance in which being motivated by and acting on that reason would be more rational than not. In this case, holding fixed Barry's intention to build a door, his knowledge of block planes, their usefulness, and their availability, Barry would be more rational to use a block plane than not. And this has everything to do with the things I'm holding fixed, the grounds of his reasons. Even fans of wide-scope oughts have little reason here to deny that a narrow-scope ought claim expressed with the same words is also true here.

Harry intends to kill his father undetected. We've already suggested that a practical conditional of the sort we're after applies to him as well as to Barry: If Harry intends to kill his father undetectably he ought to use poison. That's so because poison is an effective undetectable killing agent, and because Harry intends to do that. Harry's goal itself is of course bad, and maybe even irrational. But that fact has got very little to do with the underlying facts that make the practical conditional true. In other words, the facts which ground the conditional have everything to do with rationalizing Barry's being motivated to use poison. But they don't in fact

have much to do with the rational pressure not to intend to kill his father. Our test for reasons was to hold the grounds fixed and ask if the subject would be more or less rational to respond in one way rather than the other.<sup>16</sup> So hold fixed Harry's intention, that poison would be efficacious, and Harry's knowledge of these facts. Wouldn't Harry be more rational to use poison than not?

You might want to resist because of the badness of the end in question. Here's a principle you might be tempted to cite: When the purported grounds of a reason include features of the agent that are constitutive of irrationality, those should not be held fixed in figuring out what her reasons are. That sort of thought was what caused trouble for the "ideal rationality" formulations discussed briefly above. My psychological disorder gives me a reason to see a psychologist. But it is constitutive of irrationality that I have it. We need to take the fact that I have it into account in determining my reasons or else we wrongly conclude I have no reason to see a psychologist. If we want to resist using "bad" desires to generate reasons, we'll need some other motivation.

Reasons rationalize and explain change over time as well as rule on the rationality of total states at a time. Barry intends to make a door, but doesn't plan on or expect to use a plane.<sup>17</sup> If Barry is or becomes aware of the facts that make the conditional in question true we would expect him to change his mind about using a plane by the time using it becomes

---

<sup>16</sup>I'm not claiming that the test provides sufficient conditions for having a reason, but it surely provides evidence that the agent has a reason.

<sup>17</sup>Switching from the bare fact that he will not or does not use a plane (which actually is the consequent on the wide-scope view) to his knowledge of it doesn't seem to me to be unfair to the view. It includes that fact insofar as it is knowledge, and it actually helps make it more plausible that that fact could rationalize an action than if we left it unrecognized by Barry.

appropriate. We would cite those fact in rationalizing his coming to intend this, and we would have an explanation of why he did so. So far so good. Suppose instead that Barry predicts that he will not use a plane, and he is aware or becomes aware of the facts that make the underlying conditional true. We would not expect him to give up his plan to make a door. Or perhaps we would figure that he had given up his intention to make a door, but we would not find citing this fact – that he was not going to use a block plane – as either a justification or an explanation of his giving up his plan. We would rather look around for other explanations of his giving up the plan.

So the suggestion is that reasons must be capable of rationalizing action, and they must be capable of doing so even in people who are less than fully ideal. We should not think of reasons and the ought claims that express them as incompatible with rational failures, including the rational failures involved in having bad ends. If that's right, and my claims about what rationalizes what in the above examples are correct, we have some good reason to think that there are reasons of the sort needed for the consequents of the target practical conditionals to be true.

## Bibliography

Broome, John, 'Normative Requirements'. *Ratio* 12(4) (1999): 398-419.

\_\_\_\_\_, 'Normative Practical Reasoning', *Proceedings of the Aristotelian Society*, supplementary vol. 75 (2001): 175-193.

\_\_\_\_\_, 'Practical Reasoning', J. L. Bermudez & A. Millar (eds.), *Reason and Nature: Essays in the Theory of Rationality*, (New York: Oxford University Press, 2002).

Finlay, Stephen, 'What Ought Probably Means, And Why You Can't Detach It' (manuscript).

Geach, P. T., 'Whatever happened to deontic logic?' *Philosophia*, 11:1--12, 1981.

Greenspan, Patricia, 'Conditional Oughts and Hypothetical Imperatives', *The Journal of Philosophy* 72(10) (1975): 259-276.

Grice, Paul, *Studies in the Way of Words*, (Cambridge; Harvard University Press, 1989).

Harman, Gilbert, 'Review of Roger Wertheimer, *The Significance of Sense*', *Philosophical Review* 82, 235--39 (1973).



Hill, Thomas, 'The Hypothetical Imperative'. in his *Dignity and Practical Reason in Kant's Moral Theory*. (Ithaca, NY: Cornell University Press, 1973).

R. Johnson, 'Internal Reasons and the Conditional Fallacy', *The Philosophical Quarterly*, 49 (January 1999) pp. 53-71.

Korsgaard, Christine, 'The Normativity of Instrumental Reason', in G. Cullity & B. Gaut (eds.), *Ethics and Practical Reason*. (New York: Oxford University Press, 1997) .

Raz, Joseph, 'The Myth of Instrumental Rationality', *Journal of Ethics and Social Philosophy* 1 (2005): 2-28.

Schroeder, Mark, 'Having Reasons', *Philosophical Studies* (forthcoming).

\_\_\_\_\_ 'The Scope of Instrumental Reason', *Philosophical Perspectives*, 18 (2004): 337-364.

\_\_\_\_\_ 'Instrumental Mythology.' *Journal of Ethics and Social Philosophy*, [www.jesp.org](http://www.jesp.org), (2005).

\_\_\_\_\_ *Slaves of the Passions*, (Oxford; Oxford University Press, 2008a).

\_\_\_\_\_ 'Means-End Coherence, Stringency, and Subjective Reasons.' *Philosophical Studies*

(2008b).

Setiya, Kieran, 'Cognitivism about Instrumental Reason', *Ethics* (2007).

Shope, R.: 'The Conditional Fallacy in Contemporary Philosophy'. *Journal of Philosophy* 75 (1978): 397–413.

Sobel, David, "Explanation, Internalism, and Reasons for Action." *Social Philosophy and Policy* 18, no. 2 (Summer 2001): 218–35.

van Roojen, Mark, 'Motivational Internalism: A Somewhat Less Idealized Account', *Philosophical Quarterly*, (2000).

Wallace, R. Jay, 'Normativity, Commitment and Instrumental Reason', *Philosophers' Imprint* 1 (2001): 1-26.

Bernard Williams, 'Internal and External Reasons', in *Moral Luck*, (Cambridge: Cambridge Univ. Press, 1981) pp. 101-113.

Bernard Williams, 'Internal Reasons and the Obscurity of Blame', in *Making Sense of Humanity* (Cambridge: Cambridge University Press, 1995) pp. 35-45.