

In Defense of the Wrong Kind of Reason

Christopher Howard
University of Arizona
chrishoward@email.arizona.edu

[Abstract]

Skepticism about the ‘wrong kind’ of reasons—the view that wrong-kind reasons are reasons to want and bring about certain attitudes, but not reasons for those attitudes—is more often assumed than argued for. Jonathan Way (2012) sets out to remedy this: he argues that skeptics about, but not defenders of, wrong-kind reasons can explain a distinctive pattern of transmission among such reasons and claims that this fact lends significant support to the skeptical view. I argue that Way’s positive case for wrong-kind reason skepticism fails. I conclude with an account of what’s needed to resolve the debate between wrong-kind reason skeptics and defenders.

1. Introduction

Call any consideration that counts broadly in favor of ϕ -ing, but which doesn’t bear on whether ϕ -ing would be fitting or correct a *wrong-kind* reason (WKR) to ϕ . Assume, plausibly, that it’s correct to believe only what’s true, to intend to do only what’s worth doing, and to desire only what’s desirable. Then considerations that don’t bear on whether p but somehow favor believing p are WKRs to believe p ; considerations that don’t bear on whether A -ing is worth doing but somehow favor intending to A are WKRs to intend to A ; and considerations that don’t bear on whether x is desirable but somehow favor desiring x are WKRs to desire x .

Examples of WKRs are easy to come by. In the case of belief, think of Pascal’s Wager. In the case of intention, think of Kavka’s Toxin Puzzle (Kavka 1983). And in the case of desire, think of Roger Crisp’s evil demon that threatens to punish you unless you desire a(n) (undesirable) cup of mud (Crisp 2000). In these cases, it looks like there are good reasons for you to believe, intend, and desire respectively. But according to the standards that we’re assuming, these reasons don’t bear on whether believing, intending, or desiring would be correct. So they’re WKRs.

Contrast WKRs with *right-kind* reasons (RKR): considerations that count in favor of ϕ -ing which *do* bear on whether ϕ -ing would be correct. Examples of RKRs are even easier to come by. That it’s raining is an RKR to believe that it’s wet out.

That it would help those who need it is an RKR to intend to do some charity work. And that it's delicious is an RKR to want some chocolate ice cream. These considerations look like good reasons to believe, intend, and desire, respectively. And according to the standards that we're assuming, they bear on whether believing, intending, or desiring would be correct. So they're RKRs.

A number of philosophers are *WKR skeptics*: They hold that RKRs for ϕ -ing are the only reasons for ϕ -ing there are. According to these philosophers, putative WKRs don't give you reasons, but *incentives*, where an incentive to ϕ is a reason to want to ϕ and to try to bring it about that you ϕ , but not a reason to ϕ . Thus, according to the WKR skeptic, the fact that Crisp's demon will punish you unless you desire a cup of mud doesn't strictly speaking give you a reason to desire a cup of mud. Rather, this fact merely gives you a reason to want to desire a cup of mud and to bring this about if you can.¹

Many find the skeptical treatment of WKRs plausible, particularly in the case of belief. It's widely held that considerations irrelevant to the truth of p can't be reasons for believing p . And WKR skepticism implies this result. But some theorists—call them *WKR defenders*—remain unmoved. This, because WKR skepticism is often put forward without argument. A number of writers, including Allan Gibbard (1990), Derek Parfit (2011), and John Skorupski (2010) merely assert that WKRs for attitudes aren't reasons for those attitudes, but are instead incentives for them. But clearly we might question whether this is so. Suppose that whether Crisp's demon delivers on his threat to punish you depends on whether you *actually* come to desire a cup of mud. Then, as some have noted,² it's not at all clear that the demon's threat merely gives you reason to want and to try to desire a cup of mud, but no reason to desire a cup of mud. So WKR skepticism isn't justified on its face. The view requires independent support.

In recent work, Jonathan Way (2012) sets out to provide such support. Way's argument for WKR skepticism turns on the fact that reasons transmit: that, in many cases, there's reason for one action or attitude because there's reason for another.

¹ See, among others, Gibbard (1990); Kelly (2002); (Kolodny 2005); Owens (2000); Parfit (2011); Shah (2006, 2008); Skorupski (2010). Some of these authors endorse a *selective* WKR skepticism: they're skeptical about WKRs for belief, and neutral with respect to WKRs for desire, for example.

² See Danielsson and Olson (2007), pp. 513-14; and for similar concerns about the lack of argument for WKR skepticism, see Rabinowicz and Rønnow-Rasmussen (2004).

For example, if there's reason to want the drive to pass quickly, and listening to music would help make this so, then there's reason to want to listen to music. From the fact of transmission, Way's argument proceeds in two stages. First, he argues that WKR's transmit in a way that's distinctive of such reasons—i.e., in a way that's different from RKR's. Way suggests that this fact calls for explanation, that an account of WKR's shouldn't take the manner in which such reasons transmit as primitive. Second, he argues that while WKR skeptics can meet this explanatory task—in particular, by appealing to the way in which reasons to want and bring about attitudes transmit—WKR defenders can't. According to Way, this fact lends significant support to WKR skepticism.

In this paper, I take up Way's argument and target it at its first stage. I argue that the way in which WKR's transmit is *not* distinctive—that RKR's transmit in a parallel fashion. The upshot is that, contrary to Way, there's no special challenge for WKR defenders to explain the manner in which WKR's transmit: the way in which such reasons transmit isn't a feature of WKR's in particular but rather of reasons for attitudes *in general*. Thus WKR defenders can explain the way in which WKR's transmit by appeal to a fully general principle concerning the transmission of reasons. And this means that WKR skepticism can't gain in plausibility from its purportedly unique ability to explain why WKR's transmit in the way that they do. So Way's argument fails to provide much needed support for WKR skepticism.

2. Transmission Patterns

To start, it will be useful to consider the way in which reasons for action transmit from ends to means. So suppose there are reasons for you to take a road trip to Boston. Since there are reasons for you to take the trip, there are reasons for you to take the means to this end: to get in your car, to fill it with gasoline, to head toward Boston, and so on. Additionally, there are reasons for you to make sure that certain *enabling conditions* of your taking the trip are met: to ensure that your car is drivable, and that the roads on which you're planning to ride are clear for passage. And beyond this, there might in other cases be reasons for you to take means *constitutive* of some end: to drink some water as a way of quenching your thirst, or to lie outside as a way of catching some sun. Following Joseph Raz (2005), Way says that any action that "helps you to achieve an end" in at least any of the ways just mentioned *facilitates* that end (2012, p. 494). He goes on to claim that reasons for action transmit across facilitative connections between actions as follows:

Action Pattern. If there is a reason to *A*, then the fact that *B*-ing facilitates *A*-ing is a reason to *B*.

The Action Pattern requires some qualifications. But I won't discuss these details here. Instead, I'll assume with Way that, given the necessary qualifications, the Action Pattern holds.³

Now what's important for us is that, like reasons for action, reasons for attitudes also transmit across facilitative connections. But according to Way, WKR's transmit in a way that's distinctive of such reasons—i.e., in a way that's different from RKR's. In particular, he claims that RKR's transmit (exclusively) across facilitative connections between the *objects* of attitudes, whereas WKR's transmit across facilitative connections between *attitudes* themselves.

To bear this out, consider first the case of RKR's to desire. According to Way, such reasons transmit according to the following pattern:

Desire Pattern. If there is an RKR to desire *x*, then the fact that *y* facilitates *x* is an RKR to desire *y*.

The Desire Pattern tells us that RKR's to desire transmit across facilitative connections between the objects of desire: that when one thing helps to bring about something that there's an RKR to want, the fact that it does so is an RKR to want that thing. For example, the Desire Pattern predicts that if there's an RKR to want to drive safely, then the fact that buckling your seatbelt will help you to drive safely is an RKR to want to buckle your seatbelt.

Way doesn't provide any explanation for why RKR's to desire transmit in this way, or for why RKR's for other attitudes transmit in a parallel fashion (i.e. across facilitative connections between the objects of those attitudes). Rather he takes this fact to be common ground between WKR skeptics and defenders—a fact to which

³ Each of the patterns of transmission discussed in what follows is subject to the same set of qualifications. For a list and discussion of the relevant qualifications see Way (2012), pp. 494-96.

WKR skeptics can unproblematically appeal in arguing for WKR skepticism, without first needing to explain.⁴

However, what Way does think requires explanation is the manner in which WKRs for attitudes transmit. According to Way, such reasons transmit as follows:

Wrong Reason Pattern. If there is a WKR for attitude *A*, then the fact that attitude *B* facilitates attitude *A* is a WKR for attitude *B*.

The Wrong Reason Pattern tells us that WKRs transmit across facilitative connections between attitudes: that when one attitude facilitates an attitude for which there's a WKR, the fact that it does so is a WKR for that attitude. For example, the Wrong Reason Pattern predicts that if there's a WKR to intend today to drink a vile toxin tomorrow at noon—e.g., that an eccentric billionaire has offered you a \$1 million prize to have this intention (Kavka 1983)—then the fact that believing that you ought to drink the toxin will facilitate your intending to drink the toxin is a WKR to believe that you ought to drink the toxin.⁵

⁴ For discussion, see Rabinowicz, Rønnow-Rasmussen, and Way, “Ethics Discussion at PEA Soup: Rabinowicz and Rønnow-Rasmussen on Way,” conversation on PEA Soup, entries dated June 28, 2012 and June 29, 2012, <http://peasoup.typepad.com/peasoup/2012/06/ethics-discussions-at-pea-soup-rabinowicz-and-ronnow-rasmussen-on-way.html> [last accessed June 19, 2015]. In their comment, Rabinowicz and Rønnow-Rasmussen propose that we might explain the way in which RKR's transmit by appeal to facts about the transmission of value among objects of attitudes for which there are RKR's. This proposal looks promising (at least when it comes to RKR's for desire and intention), but Way resists it on the grounds that he wishes to defend ‘buck-passing’ accounts of value—accounts which explain value properties in terms of reasons for evaluative attitudes. Thus Way holds that facts about the transmission of value should be explained in terms of facts about the transmission of reasons, rather than the other way around, so as to avoid circularity.

⁵ This example comes from Way (2012), pp. 497-98. Way thinks that believing that you ought to drink the toxin is a means to intending to drink the toxin, since if you believed that you ought to drink the toxin, then you could form the intention to drink it on that basis (you could reason directly from the belief that you ought to drink the toxin to an intention to drink the toxin). Thus believing that you ought to drink the toxin would in this way help you to intend to drink the toxin, and so to win the billionaire's reward. This looks like a reason to believe that you ought to drink the toxin. But it's not *evidence* that you ought to drink the toxin—it's a respect in which it would be *good* to believe that you ought to drink the toxin. So it's a WKR for this belief.

Now remember: according to Way, the fact that WKR for attitudes transmit in this manner (i.e., across facilitative connections between attitudes) distinguishes them from other kinds of reasons. In particular, he claims that no pattern of transmission parallel to the Wrong Reason Pattern holds for RKR, or for reasons for attitudes in general. Way suggests that this (purported) fact calls for explanation. He goes on to argue that while WKR skeptics can meet this explanatory challenge, WKR defenders can't.⁶ According to Way, this fact provides WKR skeptics with an important advantage over WKR defenders.

But in the section that follows, I'll argue that in fact WKR skeptics have no such advantage. More specifically, I'll argue that the way in which WKR transmit is *not* distinctive of such reasons: that reasons for attitudes in general transmit according to a pattern which parallels the Wrong Reason Pattern. WKR defenders can thus explain the Wrong Reason Pattern by appeal to a fully general pattern of transmission among reasons, of which the Wrong Reason Pattern is just a special case. Thus contrary to Way, the WKR skeptic's ability to explain the Wrong Reason Pattern in fact provides her with no advantage over WKR defenders.

3. The General Pattern

To begin, consider the

General Pattern. If there is a reason for attitude *A*, then the fact that attitude *B* facilitates attitude *A* is a WKR for attitude *B*.

The General Pattern tells us that reasons for attitudes *in general*—i.e., both WKRs and RKRs alike—transmit across facilitative connections between attitudes: that when one attitude facilitates an attitude for which there's a reason, the fact that it does so is a WKR for that attitude. Thus the General Pattern combines two claims: the Wrong Reason Pattern and the

⁶ As I mentioned earlier, Way's skeptical explanation of the manner in which WKRs transmit proceeds by appeal to the way in which reasons to want and to bring about attitudes transmit. Since my claim isn't that Way's proposed skeptical explanation of the Wrong Reason Pattern fails, but rather that WKR defenders can also plausibly explain this pattern, I won't rehearse Way's proposed explanation here (but for the interested reader, see Way [2012], pp. 500-2).

Mixed Pattern. If there is an RKR for attitude *A*, then the fact that attitude *B* facilitates attitude *A* is a WKR for attitude *B*.

If the Mixed Pattern holds, then given that the Wrong Reason Pattern holds, the General Pattern holds. And if the General Pattern holds, then Way's claim that no pattern parallel to the Wrong Reason Pattern holds for reasons in general is false. So in order to maintain that the Wrong Reason Pattern is distinctive of WKRs, Way must deny the Mixed Pattern. And he does. He claims that we should reject the Mixed Pattern, and so the General Pattern.

However, before he settles on this conclusion, Way pauses to consider the following argument in favor of the Mixed Pattern:

[The Mixed Pattern] may seem plausible. For it may seem that it is good to have attitudes there is the right kind of reason for. Of course, such attitudes need not be instrumentally good. But such attitudes may be thought to be good for their own sake, insofar as they constitute a proper [i.e., correct] response to one's reasons. If so, then attitudes which facilitate attitudes there is the right kind of reason for are instrumentally good, and so attitudes there is the wrong kind of reason for. (Way 2012, p. 499)

As I indicated, Way goes on to reject this argument. But I find his reply suspect. Indeed, I think that it fails. I'll make the implications of this clear in a moment. But first, I'll rehearse Way's reply to the above argument for the Mixed Pattern, and explain how it fails.

Way's reply is this. While he accepts that it "may well be good for its own sake to have attitudes which constitute a [correct] response to one's reasons," Way says that an attitude constitutes a correct response to one's reasons only if there's *sufficient* RKR for it—i.e., only if the attitude is correct *overall* (Way 2012, p. 499). So Way thinks that having an attitude for which there's merely *some* RKR isn't necessarily correct, and so isn't necessarily good. So he thinks that attitudes that facilitate attitudes for which there are RKRs aren't necessarily instrumentally good, and thus that there aren't necessarily WKRs for such attitudes. And so Way concludes that the above argument for the Mixed Pattern fails.

Now, an axiology on which it's good to have correct attitudes is, I think, highly plausible.⁷ Other things equal, it seems good to intend to do things that are worth doing, to desire things that are desirable, and to have true beliefs. Further, it seems that the value of having such attitudes needn't have anything to do with the effects of having them. Rather, as the above argument for the Mixed Pattern suggests, having correct attitudes seems good for its own sake, or of *final* value. But according to Way, an attitude *can't* be correct, and so plausibly of any final value at all, unless there's sufficient RKR for the attitude—unless it's correct overall. And I think this is a mistake.

I assume that an attitude can be correct *to some degree*, without being correct overall—namely, in situations where there is some, but not sufficient, RKR for the attitude.⁸ So on an axiology according to which correct attitudes are of final value, I think it's natural to hold that such attitudes are *proportionally* valuable—i.e., valuable in proportion to the degree to which they're correct.⁹ Thus, given such an axiology, considerations that make an attitude to some degree *correct*—i.e., RKRs for the attitude—plausibly count as, or “add up to”, considerations that make having that attitude to some degree *finally valuable*. Consequently, an axiology according to which correct attitudes are of final value plausibly yields the following result:

RKR-Value Link (RVL). A consideration, *p*, is an RKR for an attitude, ϕ , only if *p* makes it the case that ϕ -ing would be finally valuable to some degree.

RVL is, I think, an attractive axiological hypothesis. It predicts, e.g., that believing *p* is *pro tanto* good when there are RKRs to believe *p*; that intending to *A* is *pro tanto* good when there are RKRs to intend to *A*; and that desiring *x* is *pro tanto* good when there are RKRs to desire *x*.¹⁰ And these look like compelling axiological results. But

⁷ Cf. Hurka (2001), esp. ch. 1, and Sylvan (2012).

⁸ Since I take ‘correct’ to be synonymous with ‘fitting’, I think this assumption should be uncontroversial. Heathwood (2008) points to the fact that ‘fitting’ is a gradable adjective (p. 58). And so too do Chappell (2012) and Zimmerman (2011), among others. So at the very least, my assumption that ‘correct’ (viz., ‘fitting’) is a gradable adjective isn't without precedent.

⁹ For an extended defense of this proportionality thesis, see Hurka (2001), esp. ch. 2.

¹⁰ Note that RVL does not entail that if there's an RKR for an attitude, then having that attitude is good *overall*. All RVL says is that whenever there's an RKR for an attitude, having

what's important about RVL, specifically in the present context, is that it provides an argument for the Mixed Pattern.

The argument from RVL to the Mixed Pattern is straightforward. RVL implies that, necessarily, if there's an RKR for an attitude, then having that attitude is finally valuable to some degree. So if an attitude facilitates an attitude that there's an RKR for, then having that attitude is necessarily instrumentally valuable to some degree. And if an attitude is instrumentally valuable (to any degree), then there's necessarily a WKR for that attitude. So, given RVL, there are necessarily WKRs for attitudes that facilitate attitudes for which there are RKRs. And this means that RKRs transmit across facilitative connections between attitudes according to the Mixed Pattern. Thus, given an independently plausible, substantive axiological hypothesis according to which RKRs for attitudes are considerations that make having those attitudes (to some degree) finally valuable, the Mixed Pattern holds. So I think we should say that this pattern holds.

However, if the Mixed Pattern holds, then given that the Wrong Reason Pattern holds, the General Pattern holds. And thus, Way's claim that no pattern parallel to the Wrong Reason Pattern holds for reasons for attitudes in general is false: the Wrong Reason Pattern isn't distinctive of WKRs. The upshot is that, contrary to Way, there's no special challenge for WKR defenders to explain the Wrong Reason Pattern. For WKR defenders can plausibly explain this pattern by appeal to the General Pattern. In particular, they can say that WKRs transmit according to the Wrong Reason Pattern because WKRs for attitudes are *reasons for those attitudes*, and all reasons for attitudes transmit according to the General Pattern—a pattern of which the Wrong Reason Pattern is just a special case. Thus, WKR defenders have available to them a plausible explanation of the Wrong Reason Pattern. And this means that the WKR skeptic's ability to explain this pattern provides her with no advantage over WKR defenders.

4. An Alternative Formulation?

I've argued that Way's recent positive case for WKR skepticism falters at its first step—that the way in which WKRs transmit isn't distinctive of such reasons—and that this fact allows WKR defenders to plausibly explain the pattern according to

that attitude is good *to some degree*. Indeed, RVL is consistent with the possibility of an attitude's being bad overall (either instrumentally or for its own sake), even when there are RKRs for that attitude.

which WKR transmit (the Wrong Reason Pattern) by appeal to a fully general pattern of transmission among reasons (the General Pattern).

However, a natural reformulation of Way's argument is in the offing, and so is worth considering. So in this last section, I'll conclude by considering it.

The transmission profiles of WKRs and RKR might differ in two ways: WKRs might transmit in a way that RKR don't or WKRs might fail to transmit in a way that RKR do. I've argued that the transmission profiles of WKRs and RKR don't differ in the first way: that RKR transmit in the way that WKRs do, viz., across facilitative connections between attitudes. However, I do think they differ in the second way: that RKR transmit in a way that WKRs don't, viz., across facilitative connections between the objects of attitudes.¹¹ To see this, consider the

Wrong Desire Pattern. If there is a WKR to desire x , then the fact that y facilitates x is a WKR to desire y .

The Wrong Desire Pattern is the Desire Pattern with "RKR" replaced by "WKR". It tells us that WKRs for desire transmit across facilitative connections between the objects of desire: that when one thing helps to bring about something that there's a WKR to want, the fact that it does so is a WKR to want that thing. But this is false. If a demon threatens to punish you unless you want a cup of mud, then the fact that acquiring a cup and finding some mud facilitates getting a cup of mud might be an RKR to want to do this—if there's anything desirable about having a cup of mud—but it's not a WKR for this desire.¹² So the Wrong Desire Pattern doesn't hold.

To generalize: whereas RKR transmit across facilitative connections between the objects of attitudes—i.e., in accordance with patterns like the Desire Pattern—WKRs don't. Rather WKRs seem to transmit *only* across facilitative connections between attitudes themselves—i.e., in accordance with the Wrong Reason and General Patterns. So if WKR skeptics could explain *this* difference in the ways in which WKRs and RKR transmit—and if WKR defenders could provide no similarly plausible explanation of the difference—this would lend some support to WKR skepticism. And thus, a natural reformulation of Way's argument is in the offing,

¹¹ Way (2012, p. 497) notes this difference in the transmission profiles of WKRs and RKR in passing, but doesn't make any hay out of it.

¹² I borrow this example from Way (2012, p. 497).

albeit one which takes as a premise a different difference in the ways in which WKR and RKR transmit.

And in fact, WKR skeptics can plausibly explain the difference in question—they can plausibly explain why, unlike RKR, WKRs seem to fail to transmit according to patterns like the Desire Pattern. In particular they can do so by appealing to the ways in which RKR to want attitudes—i.e., RKR for higher-order desires—fail to transmit. To see this, consider the

Desire to Desire Pattern. If there is an RKR to want to want x , then the fact that y facilitates x is an RKR to want to want y .

The Desire to Desire Pattern is the Desire Pattern with “RKR” replaced by “RKR to want”. It tells us that RKR for higher-order desire transmit across facilitative connections between the objects of lower-order desire: that when one thing helps to bring about something that there’s an RKR to want to want, the fact that it does so is an RKR to want to want that thing. But this is false. If there’s an RKR to want to want a cup of mud, then the fact that acquiring a cup and finding some mud facilitates getting a cup of mud might be an RKR to want to do this—if there’s anything desirable about having a cup of mud—but it’s not an RKR to want to want to do this.

To generalize: just as WKRs seem to fail to transmit according to patterns like the Desire Pattern, so too do RKR for higher-order desires. And thus, WKR skeptics can plausibly explain the former fact by appeal to the latter: they can say that WKRs seem to fail to transmit according to patterns like the Desire Pattern because WKRs are in fact just RKR for higher-order desires, and RKR for higher-order desires fail to transmit according to such patterns.

So WKR skeptics can plausibly explain why, unlike RKR, WKRs seem to fail to transmit according to patterns like the Desire Pattern. So if WKR defenders can’t give a similarly plausible explanation of this difference, this would lend support to WKR skepticism.

However, WKR defenders can give a similarly plausible explanation of the difference. In fact, they can do so in much the manner that WKR skeptics can. Let me explain.

According to WKR skeptics, WKRs for attitudes are RKR to want those attitudes, but not reasons for those attitudes. A number of WKR defenders happily accept the first, positive part of this claim: that something is a WKR for an attitude

just in case it's an RKR to want that attitude. What they deny is the second, negative part: that RKR to want attitudes aren't also reasons for those attitudes (Danielson and Olson 2007, Louise 2009, Markovits 2010, Rabinowicz and Rønnow-Rasmussen 2004). And for WKR defenders who accept this view, explaining the difference with which we're concerned is simple: they can say that WKRs fail to transmit according to patterns like the Desire Pattern because WKRs are *equivalent* to RKR for higher-order desires, and RKR for higher-order desires fail to transmit according to such patterns.

The plausibility of this explanation on the part of the WKR defender turns on the plausibility of her claim that RKR to want some attitude are also reasons for that attitude. But fortunately—and contrary to WKR skeptics—this claim is plausible. Consider the normative thesis that p is an RKR to want to do some act only if p is a reason to do that act. This thesis is attractive and widely accepted.¹³ Why, then, shouldn't we think that a parallel thesis holds in the case of reasons for attitudes? If RKR to want to do some act are reasons to do that act, then why shouldn't RKR to want some attitude be reasons for that attitude? True, WKR defenders who accept this thesis ultimately owe an explanation of its truth.¹⁴ However, given the plausibility of its parallel in the case of reasons to act, WKR skeptics similarly owe an explanation of its negation. So, as it stands, WKR skeptics have no more of a claim to the thesis that RKR to want attitudes *aren't* also reasons for those attitudes than WKR defenders do to the thesis that they *are*. Settling this matter is, I believe, the way forward in the debate between WKR skeptics and WKR defenders.

References

- Chappell, R. "Fittingness: The Sole Normative Primitive." *Philosophical Quarterly* 62 (2012): 684-704.
- Clarke, R. "Some Theses on Desert." *Philosophical Explorations* 16 (2013): 153-64.
- Crisp, R. "Review of John Kupperman, Value...and What Follows." *Philosophy* 75 (2000): 458-62.

¹³ What's more controversial is the thesis that something is a reason to do some act only if it's an RKR to want to do that act (though see Portmore [2011] for a recent defense).

¹⁴ On this, see Way (2012). I offer such an explanation in Howard (unpublished). But for reasons of space, I can't discuss the details here.

- Danielsson, S. and J. Olson. "Brentano and the Buck-Passers." *Mind* 116 (2007): 511-22.
- Gibbard, A. *Wise Choices, Apt Feelings*. Cambridge, MA: Harvard University Press, 1990.
- Heathwood, C. "Fitting Attitudes and Welfare." In *Oxford Studies in Metaethics*, vol. 3, edited by R. Shafer-Landau, 47-73. Oxford: Oxford University Press, 2008.
- Howard, C. "The Fundamentality of Fit." Unpublished manuscript.
- Hurka, T. *Virtue, Vice, and Value*. Oxford: Oxford University Press, 2001.
- Kavka, G. "The Toxin Puzzle." *Analysis* 43 (1983): 33-36.
- Kelly, T. "The Rationality of Belief and Some Other Propositional Attitudes." *Philosophical Studies* 110 (2002): 163-96.
- Louise, J. "Correct Responses and the Priority of the Normative." *Ethical Theory and Moral Practice* 12 (2009): 345-64.
- Markovits, J. "Internal Reasons and the Motivating Intuition." In *New Waves in Metaethics*, vol. 6, edited by R. Shafer-Landau, 255-79. Oxford: Oxford University Press, 2011.
- Niko, K. "Why Be Rational?" *Mind* 114 (2005): 509-63.
- Owens, D. *Reasons without Freedom*. Oxford: Routledge, 2000.
- Parfit, D. *On What Matters*. Oxford: Oxford University Press, 2011.
- Portmore, D. "The Teleological Conception of Practical Reasons." *Mind* 120 (2011): 117-53.
- Rabinowicz, W. and T. Rønnow-Rasmussen. "The Strike of the Demon: On Fitting Pro-Attitudes and Value." *Ethics* 114 (2004): 391-423.
- Raz, J. "The Myth of Instrumental Rationality." *Journal of Ethics and Social Philosophy* 1 (2005): 1-28.
- Shah, N. "A New Argument for Evidentialism." *Philosophical Quarterly* 56 (2006): 481-98.
- . "How Action Governs Intention." *Philosopher's Imprint* 8 (2008): 1-19.
- Skorupski, J. *The Domain of Reasons*. Oxford: Oxford University Press, 2010.
- Sylvan, K. "Truth Monism without Teleology." *Thought* 1 (2012): 161-69.
- Thomas, K. "The Rationality of Belief and Some Other Propositional Attitudes." *Philosophical Studies* 90 (2002): 163-96.
- Way, J. "Transmission and the Wrong Kind of Reason." *Ethics* 122 (2012): 489-515.
- Zimmerman, M. "Partiality and Intrinsic Value." *Mind* 120 (2011): 447-83.

