

The Fundamentality of Fit

Christopher Howard
UNC-Chapel Hill

forthcoming in *Oxford Studies in Metaethics*, vol. 14
Winner of the Marc Sanders Prize in Metaethics

1. Introduction

The recently influential “reasons-first” approach to normativity says that reasons are the fundamental elements of the normative domain, and that all other normative facts, properties, and relations can be analyzed or accounted for in terms of the reason relation (Parfit 2011, Scanlon 1998, Schroeder 2007). A central alternative—the “value-first” approach—says that value or goodness occupies the fundamental position in the ontology of normativity, and that the rest of the normative can be accounted for in terms of the property of being good (Moore 1903). In the first part of this paper, I argue that both of these approaches should be rejected. An ontology of normativity should accommodate our commonsense normative intuitions. In particular, it should accommodate the existence of all of the normative reasons that, intuitively, there are. But both the reasons- and value-first approaches fail in this regard: neither can be reconciled with the existence of a certain *kind* of reason.

In the second part of the paper, I advance an alternative ontology of normativity—one that takes *fittingness* as fundamental. The normative relation of fittingness is the relation in which a response stands to an object when the object merits—or is *worthy of*—that response. Call the approach that says that fittingness is the fundamental feature of normativity, and that the rest of the normative can be accounted for in terms of fittingness, the “fittingness-first” approach to normativity.

The fittingness-first approach has a rich history. Versions of the ontology were originally suggested by Franz Brentano (1889/2009) and A.C. Ewing (1948), and have recently been proposed by more contemporary authors (Chappell 2012, McHugh and Way 2016). I’ll argue, however, that existing versions of the fittingness-first approach to normativity suffer the same problem as the reasons- and value-first approaches: none can accommodate all the normative reasons there are.

Because my version of the fittingness-first approach similarly accounts for all of normativity in terms of just a single normative relation, it’s no less parsimonious than the reasons-first approach, the value-first approach, or existing versions of the

fittingness-first approach. And since, as I'll argue, my fittingness-first view can plausibly accommodate and explain the existence of all the normative reasons there are, it provides a substantively superior and so overall more attractive ontology of normativity. Thus, as compared to the reasons- and value-first approaches, as well as existing versions of the fittingness-first approach, my version of the fittingness-first approach ought to be preferred. This is the central thesis of the paper.

Two preliminary remarks are in order. First, the “analyses” or “accounts” of various normative items that I'll be interested in throughout the paper are intended as metaphysical accounts of the natures or essences of those items, as opposed to linguistic or conceptual analyses.¹ Second, I should be clear that my fittingness-first view claims only that fittingness is *normatively* fundamental—i.e., fundamental relative to the rest of the normative domain. This leaves open the possibility that fittingness isn't fundamental *simpliciter*—that the relation can be fully accounted for in naturalistic or descriptive terms. I'll remain agnostic about this possibility here.

2. Two Kinds of Reasons

Normative reasons *count in favor* of what they're reasons for.² So when it comes to attitudes like belief and admiration, something is a reason when it counts in favor of the attitude. But it's familiar that a fact can count in favor of an attitude in two very different ways.³

One way in which a fact can count in favor of an attitude is if it explains why the attitude would be fitting to its object. For example, the fact that Sharon spends a great deal of her time doing charity work is a fact that counts in favor of admiring Sharon, since it's a fact that explains why she's *admirable*, and so worthy

¹ For discussion of the notion of a metaphysical account of a property/relation, see, e.g., Fine (1994), Jackson (1997), Rosen (2015b), Schroeder (2007: 61-83) and Wedgwood (2007: 135-52).

² For the rest of this paper, unless I indicate otherwise, 'reason' means 'normative reason'.

³ This distinction is commonly analyzed in terms of 'object-given' versus 'state-given' reasons, and plausibly tracks Raz's (2009) distinction between 'adaptive' and 'practical' reasons for attitudes. I avoid the latter terminology, since I want to steer clear of Raz's idiosyncratic theoretical commitments concerning the 'adaptive'/'practical' distinction. And I reject the former analysis because, as is now well-recognized, it's extensionally inadequate to the phenomenon (see note 18 for further discussion). A plausible hypothesis is that the relevant distinction is coextensive with the familiar distinction between reasons of the 'right' and 'wrong' kind. But this requires defense, since the question of how to distinguish right- from wrong-kind reasons is a matter of debate. Later, I'll offer some support for this hypothesis.

of—or *fit* for—admiration. Similarly, the fact that it’s raining counts in favor of believing that it’s wet out, since it’s a fact that explains why this proposition is *credible*, and so belief-worthy, or fitting to believe.⁴

A second way in which a fact can count in favor of an attitude is if it explains why the attitude would be somehow valuable or good. For example, the fact that a deplorable dictator will order your execution unless you admire him is a fact that counts in favor of your admiring the dictator, since it’s a fact that explains why your admiring him would be good. Likewise, the fact that Sam will survive her illness only if she believes that she will is a fact that counts in favor of Sam’s believing that she’ll survive, since it’s a fact that explains why it would be good if Sam had this belief.

So, intuitively, if a fact explains why some attitude would be fitting, that fact counts in favor of that attitude. And, intuitively, if a fact explains why some attitude would be good to have, that fact counts in favor of that attitude. So if reasons for attitudes are facts that count in favor of those attitudes, then reasons for attitudes seem to come in two quite different kinds: those that explain why an attitude would be fitting and those that explain why an attitude would be good. Call reasons of the former kind *fit-related reasons*. Call reasons of the latter kind *value-related reasons*.⁵

A normative view that says that both fit- and value-related reasons are genuine reasons is highly plausible. Reasons for an attitude are facts that are relevant to determining whether you ought to have that attitude, all-things-considered. And both fit- and value-related reasons seem entirely relevant to determining what attitudes you ought, all-things-considered, to have. It’s an attractive normative hypothesis that we ought normally to have fitting attitudes: that, typically, we ought to believe only what’s credible, admire only what’s admirable, and intend to do only what’s worth doing. But in circumstances in which your having an unfitting attitude is necessary to bring about a great good (or to prevent a great evil) it seems that, all-things-considered, you ought to have that attitude, despite its being unfitting. Any contrary normative view would, it seems, stray too far from common sense.⁶

⁴ The issue of what properties comprise *credibility* is, in principle, just as much a matter of substantive normative dispute as the issue of what properties comprise, e.g., *admirability*. There is wide agreement that credibility—what merits belief—is truth. But disagreement is possible (if not actual), and so I’ll remain neutral with respect to this issue as far as possible.

⁵ The distinction between ‘fit-related’ and ‘value-related’ reasons generalizes to apply to a wide range of attitudes, including, e.g., desire, resentment, intention, trust, blame, and fear.

⁶ This isn’t to say that a normative view on which both fit- and value-related reasons are genuine reasons is uncontroversial. While the view that fit-related reasons are genuine rea-

But neither the reasons- nor the value-first approach to normativity can accommodate a normative view on which both fit- and value-related reasons are genuine reasons. The following two sections explain why.

3. The Reasons-First Approach and Value-Related Reasons

Reasons-firsters can't accommodate a view on which both fit- and value-related reasons are genuine reasons, because if value-related reasons are genuine reasons, their account of value is false. Consider the reasons-firster's account of value:

Value-in-terms-of-Reasons (V_R) For p to be a respect in which x is good is for p to be a reason to desire x .⁷

Since V_R tells us *what it is* for p to be a respect in which x is good, it entails that p is a respect in which x is good just in case p is a reason to desire x . And this entails that p can't be a reason to desire x unless p is a respect in which x is good. But not all facts that make desiring x good—i.e., value-related reasons to desire x —are also respects in which x is good. For example, the fact that a demon will kill you unless you desire a cup of mud is a fact that makes your *desiring* a cup of mud good, but it's not also a respect in which *a cup of mud* is good.⁸ And we needn't resort to such fantastical cases. For instance, the fact that wanting to quit smoking will help you to quit is a fact that makes this *desire* good to have, but it's not also a respect in which *your quitting smoking* is good. So, if value-related reasons are genuine reasons, V_R is false.

sons is among the most widely shared in contemporary normative theory, the issue of whether value-related reasons are genuine reasons for attitudes is contentious. In the following section, I discuss the prospects for skepticism about value-related reasons in some detail.

⁷ V_R is in the first instance an account of both final and instrumental value (though, according to some, there are prospects for generalizing the account to other kinds of value, e.g., to goodness-for and the varieties of attributive goodness [see Schroeder 2010: 25-55; Skorupski 2010]). Reasons-firsters will vary on the exact formulation. For example, some might want to account for value in terms of reasons for a *range* of pro-responses, including, but not limited to, that of desire; e.g., admiration, choice, wish, taking pleasure in, being glad that, and so on. And some reasons-firsters might want to specify whether a fact gets to be a respect in which x is good only if it's a reason for *everyone* to desire x , or whether it suffices for a fact's being a respect in which x is good that it be a reason for only *some* people to desire x . But for our purposes, these details won't matter. My argument can be extended to all variants of V_R .

⁸ This example famously comes from Crisp (2000).

The observation that V_R is incompatible with the existence of value-related reasons is by no means new. Famously, this is known as the “Wrong Kind of Reason” problem (WKR problem), because if value-related reasons are genuine reasons for attitudes, then these reasons are of the ‘wrong kind’ to figure in V_R .⁹

Reasons-firsters have tried to answer the WKR problem in either of two ways. Some opt to reject the counterexamples that constitute the problem—they simply deny that value-related reasons are genuine reasons for attitudes. Others concede the relevant counterexamples, try to draw a distinction between reasons that are of the ‘right’ and ‘wrong’ kind to figure in V_R , and revise the account such that it makes reference only to reasons of the ‘right’ kind. But as I’ll now explain, both of these responses are highly difficult to defend, especially for reasons-firsters.

Consider first those who try to answer the WKR problem by rejecting the counterexamples that constitute it—i.e., by denying that value-related reasons are genuine reasons for attitudes. Call such theorists *WKR skeptics*. According to WKR skeptics, value-related reasons aren’t reasons, but *incentives*, where an incentive for an attitude A is a reason to want to have A and to try to bring it about that you have A , but not a reason to have A .¹⁰ Thus, according to the WKR skeptic, the fact that a demon will kill 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 (Skorupski 2010, Way 2012).

Now perhaps we shouldn’t go so far as to call WKR skepticism a “crazy normative position” (D’Arms and Jacobson 2000b: n.61). But the suggestion that the threat of your own death doesn’t even give you *some* reason to desire the mud does seem to me deeply implausible. And the WKR skeptic’s attempt to save the phenomenon by claiming that, although there is no reason for you to desire the mud, there are reasons for you to want and to try to desire the mud just doesn’t suffice. For one, this view entails that, in the case of the demon, if you refrain from desiring a cup of mud while simultaneously wanting and trying to desire a cup of mud, then your mental life will be exactly as it should be, all-things-considered. However, as a number of authors recognize, this seems intuitively unacceptable (ibid., Rabinowicz and Rønnow-Rasmussen 2004, Rosen 2015a). For another, as Sven Danielsson and Jonas Olson (2007) rightly observe, once we specify that whether the demon delivers on his threat depends on whether you *actually* come to desire the mud, it becomes

⁹ See D’Arms and Jacobson (2000a) and Rabinowicz and Rønnow-Rasmussen (2004).

¹⁰ I borrow this characterization of “incentives” from Way (2012: 492.)

even harder to see why 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 (p. 514).

But even setting these worries aside, the real trouble is, as Danielsson and Olson put it, “that it’s difficult to find an independent rationale for why we should distinguish in the relevant way between reasons for an attitude and [reasons to want] that attitude” (ibid.: 513). To be sure, facts that make it good to have an attitude—i.e., value-related reasons for the attitude—are reasons to want that attitude: such facts make the attitude *fitting to want*, and are thus fit-related reasons to want it.¹¹ But why shouldn’t facts that are fit-related reasons to want some attitude also be reasons *for* that attitude? Consider the normative thesis that p is a fit-related reason to want to do some act only if p is a reason to do that act. This thesis is attractive and widely accepted.¹² Why, then, shouldn’t we think that a parallel thesis holds in the case of reasons for attitudes? If fit-related reasons to want to do some act are reasons to do that act, then why shouldn’t fit-related reasons to want some attitude be reasons for that attitude? And indeed, many theorists take it as a datum that fit-related reasons to want an attitude count as, or “add up to”, reasons for that attitude (ibid., Markovits 2011, Rabinowicz and Rønnow-Rasmussen 2004).¹³ So WKR skepticism isn’t justified on its face. And so in order to provide reasons-firsters with an adequate answer to the WKR problem, WKR skepticism needs independent support.

But providing WKR skepticism with the independent support it needs is no easy task, especially for reasons-firsters. Perhaps the most prominent line of argument in favor of WKR skepticism claims that facts like that of the demon’s threat—i.e., value-related reasons—aren’t genuine reasons for attitudes, on the grounds that such facts fail a ‘response-condition’ on reasons, according to which a fact p can be a reason to A only if it’s possible to A for the reason that p (Kelly 2002, Shah 2006, Parfit 2011, Gibbons 2013). But this line of argument is unavailable to reasons-firsters, since these theorists are committed to there being reasons that *violate* a response-

¹¹ This proposed connection between facts that make an attitude good to have and facts that make that attitude fitting to want is explicitly motivated and defended below.

¹² What’s more controversial is the thesis that something is a reason to do some act only if it’s a fit-related reason to want to do that act (though see Portmore, 2011 for a defense).

¹³ While the thesis that fit-related reasons to want some attitude are reasons for that attitude is plausible, those who accept it (and its parallel in the case of reasons to act) ultimately owe some explanation of *why* these theses are true—they shouldn’t be taken as brute. So far, this explanatory task has been neglected in the literature. But later (in sect. 5.2), I satisfy it.

condition, insofar as they accept V_R . To see this, suppose that a deplorable dictator threatens to start a third world war unless you admire him without also wanting to admire him.¹⁴ Since admiring the dictator would be good insofar as it would prevent another world war, V_R entails that the fact that it would prevent another world war is a reason to want to admire the dictator. However, if you form a desire to admire the dictator, then the reason that there is for you to desire to admire him will disappear, since, by stipulation of the case, if you desire to admire the dictator, he'll start another world war anyway. So V_R predicts that there's a reason to want to admire the dictator that can't survive being responded to.¹⁵ And so defenders of V_R are committed to there being reasons that violate a response-condition.¹⁶ So the WKR skeptic's appeal to a response-condition in order to explain why value-related reasons aren't genuine reasons doesn't seem promising, especially for reasons-firsters.¹⁷

So what's perhaps the most prominent line of argument for WKR skepticism doesn't suffice to provide the view with the independent support that it needs in order to offer reasons-firsters an adequate answer to the WKR problem. And so WKR skepticism is, I think, untenable as a response to this problem.

¹⁴ Reisner (2009: 271) deploys cases with a similar structure in the course of raising a different, albeit related, objection to WKR skepticism. He calls these cases of "blocked ascent".

¹⁵ This case is parallel in structure to Schroeder's (2007) well-known surprise party case, which Schroeder takes to be a counterexample to a response-condition on reasons (the case also appears in Markovits, 2011). In Schroeder's case, we're asked to consider Nate, who loves successful surprise parties thrown in his honor, but hates unsuccessful surprise parties. If there is an unsuspected surprise party waiting for Nate next door, then that fact is plausibly a reason for Nate to go next door. But, as Schroeder observes, it's simply impossible for Nate to go next door for this reason: "for as soon as you tell him about it, it will go away" (p. 165).

¹⁶ Way (2012: 512f.) makes a similar point in the context of providing an argument for WKR skepticism that doesn't rely on the truth of a response-condition. For reasons of space, I won't discuss Way's argument for WKR skepticism here, but see Howard 2016a for my response. Also, it's worth noting that, as McHugh and Way (2016) acknowledge, Way (2012) provides only an argument *that* value-related reasons aren't genuine reasons; he doesn't attempt to give an explanation of *why* such reasons aren't reasons. And in order to make an adequate case for WKR skepticism, particularly as an answer to the WKR problem, it seems WKR skeptics will need to succeed in *both* of these tasks. Establishing and explaining the truth of a response-condition on reasons would plausibly allow WKR skeptics to do just this. However, as I've just argued, the prospects for this strategy seem dim for reasons-firsters.

¹⁷ For a very different sort of argument against a response-condition, see Howard (2016b).

Consider next those who respond to the WKR problem by trying to draw a distinction between reasons that are of the ‘right’ and ‘wrong’ kind to figure in V_R . If such a distinction can be drawn, reasons-firsters can revise V_R . They can say that:

Value-in-terms-of-RKRs (V_{RKR}) For p to be a respect in which x is good is for p to be a *right-kind* reason to desire x .

Since this second strategy doesn’t require reasons-firsters to accept and defend the highly committing hypothesis that value-related reasons aren’t genuine reasons for attitudes, it’s in this respect more attractive than WKR skepticism. Nonetheless, this second strategy is also highly difficult to defend, particularly for reasons-firsters.

Perhaps the most obvious way to draw the right-/wrong-kind reason distinction is in terms of *goodness*: Right-kind reasons to want x are reasons to want x that are also facts about x that make x good. Wrong-kind reasons to want x are reasons to want x that are not also facts about x that make x good. But of course, reasons-firsters can’t draw the distinction in this way, on pain of circularity.

A second, and I think tellingly natural, way of drawing the right-/wrong-kind reason distinction is in terms of *fittingness*. Right-kind reasons to want x are reasons to want x that are also facts about x that make x fitting to want. Wrong-kind reasons to want x are reasons to want x that aren’t also facts about x that make x fitting to want. This way of drawing the distinction is attractive because, intuitively, facts about x that make x good are also facts that make x *worthy of being desired*—i.e., fitting to want. So if right-kind reasons to want x are facts about x that make x good, then, intuitively, they’re also facts that make x fitting to want. And because wrong-kind reasons to want x *aren’t* also facts about x that make x good, they’re not also facts that make x fitting to want. So, I think it’s highly natural to think that the right-/wrong-kind reason distinction can be drawn in terms of fittingness.

But reasons-firsters *can’t* draw the right-/wrong-kind reason distinction in terms of fittingness, or in terms of any normative property other than that of being a reason. This is because, according to the reasons-first approach, all normative properties ultimately need to be accounted for in terms of reasons. So reasons-firsters can draw the right-/wrong-kind reason distinction in terms of some normative property *other* than that of being a reason only if the relevant property can itself be accounted for in terms of reasons. But accounting for *any* normative property used to distinguish right- from wrong-kind reasons in terms of reasons would itself require a solu-

tion to the WKR problem. So reasons-firsters have *very* few resources with which to distinguish right- from wrong-kind reasons: they must draw this distinction in terms of reasons and non-normative properties alone—an exceedingly tall order. Although several authors have tried to offer an account of the distinction that meets this constraint, no such account has gained wide acceptance. Indeed, all such accounts of the distinction seem either subject to counterexample or to turn out to tacitly presuppose the property of being good, thus making V_R viciously circular.¹⁸

I think we should conclude that neither of the reasons-firster’s strategies for responding to the WKR problem can plausibly succeed. The upshot: reasons-firsters can neither plausibly deny nor reconcile their approach to normativity with a normative view that says that value-related reasons are genuine reasons for attitudes.

¹⁸ Three accounts of the right-/wrong-kind reason distinction in terms of reasons and non-normative properties alone have been advanced in the literature. The most famous of these is the ‘object-given/state-given’ theory, put forward by Stratton-Lake (2005), Piller (2006), and Parfit (2011). However, as I noted earlier (see note 3), this account is extensionally inadequate to the phenomenon. According to the object-given/state-given theory, right-kind reasons are equivalent to object-given reasons and wrong-kind reasons are equivalent to state-given reasons. But, as is now familiar, there are object-given wrong-kind reasons (Rabinowicz and Rønnow-Rasmussen, 2004) as well as state-given right-kind reasons (Schroeder, 2012). So the object-given/state-given theory is extensionally inadequate. A second account is that of Schroeder (2010), according to which right-kind reasons with respect to *A*-ing are those reasons shared by necessarily everyone engaged in *A*-ing, just because they’re so engaged. Schroeder’s account requires Schroeder to establish two claims: (1) that there’s a set of reasons shared by necessarily everyone engaged in *A*-ing, just because they’re so engaged and (2) that this shared set of reasons is coextensive with the set of right-kind reasons for *A*-ing. Schroeder offers two separate strategies for establishing these claims. But Sharadin (2013, 2015) convincingly argues that both of these strategies fail. And McHugh and Way (2016) raise further worries for Schroeder’s account. And Schroeder himself has abandoned it (2012). A third account comes from D’Arms and Jacobson (2000a). According to it, right- but not wrong-kind reasons to *A* are reasons to *A* which bear on the truth of *A*’s content (see also Rosen 2015a). This account faces two related problems: (1) it requires us to take on substantial and contentious commitments concerning the contents of various types of attitude in order to ensure the account’s extensional correctness (McHugh and Way 2016; Schroeder 2010); (2) these commitments would seem to render a reasons-based account of value circular, on the plausible assumption that pro-responses like desire have normative contents—for example, if desiring *x* involves seeing *x* as somehow good (Ross 1939: 278-79).

4. The Value-First Approach and Fit-Related Reasons

Unlike reasons-firsters, value-firsters *can* accommodate value-related reasons. But they can't accommodate fit-related reasons, since if fit-related reasons are genuine reasons, their account of reasons is false. On the value-firster's account of reasons:

Reasons-in-terms-of-Value (R_V) For p to be a reason to ϕ (where ϕ ranges over actions and attitudes) is for p to explain why ϕ -ing would be good in some respect and to some degree.¹⁹

Since R_V tells us what it is for p to be a reason to ϕ , it entails that p is a reason to ϕ just in case p explains why ϕ -ing would be somehow good. But not all facts that make an attitude fitting also explain why it would be good to have that attitude. For example, the fact that there are eight specks of dust on my desk is a fact that makes it fitting to believe that the number of specks of dust on my desk is even, but this fact doesn't also explain why it would be good if I had this belief. For there is plausibly no value at all in my having even a true and warranted belief about the number of specks of dust on my desk: worlds in which I have such a belief are no better than those in which I don't.²⁰ And there are still further examples, using other attitudes. For instance, the fact that your colleague just received a promotion is a fact that makes it fitting to envy her, but it's not also a fact that explains why your envying her would be good. And the fact that your friend insists on talking to you (incessantly) throughout the movie is a fact that makes it fitting to be annoyed with him, but this fact doesn't also explain why it would be good if you were annoyed with your friend. So if fit-related reasons are genuine reasons, R_V is false.

Like the observation that V_R is incompatible with the existence of value-related reasons, the observation that R_V is incompatible with the existence of fit-related reasons isn't new. Though the problem hasn't received nearly as much attention as the WKR problem, it gets some discussion in recent work by John Brunero (forthcoming) and Jonathan Way (2013). These authors don't give the problem a

¹⁹ The phrase "in some respect and to some degree" is borrowed from Raz (1999: 23), who suggests a similar view. See also Finlay (2006: 7-8; and esp. 2012) for an extended defense.

²⁰ Even most 'epistemic teleologists' are willing to admit as much. See, e.g., Goldman (1999) and Alston (2005). For an exception, see Lynch (2004), who insists that a true belief about the number of dust-specks on my desk *would* be good, insofar as it's true. I find this implausible. But nothing crucial hinges on this—I go on to provide further counterexamples to R_V .

name, but for parity's sake, we can call it the *Right Kind of Reason problem* (RKR problem) since, as I've already suggested, the reasons that cause trouble for R_V —fit-related reasons—are very plausibly the *right kind* of reasons to figure in V_R .

Now, as far as I know, no value-firster has offered a systematic response to the RKR problem in print.²¹ But there are two clear options. First, value-firsters might be “RKR skeptics”, denying that fit-related reasons are genuine reasons for attitudes. By contrast with WKR skepticism, however, RKR skepticism has no following to speak of. The view that fit-related reasons are genuine reasons is among the most widely shared in contemporary normative theory. And for good reason. It's highly plausible that the facts that make something *admirable*, and so fitting to admire, are reasons to admire it; that the facts that make something *envious*, and so fitting to envy, are reasons to envy it; that the facts that make something *credible*, and so fitting to believe, are reasons to believe it; and that the facts that make something *fearsome*, and so fitting to fear, are reasons to fear it. So I propose to set RKR skepticism aside. If I'm wrong to reject RKR skepticism—if, contrary to what I expect, a view which denies that fit-related reasons are genuine reasons can be made plausible—then the discussion to follow would have to be modified accordingly.

Second, value-firsters might adopt a theory of value according to which facts that make an attitude fitting *are* also facts that explain why having the attitude would be good in some respect and to some degree. How might this go? Some philosophers find plausible a theory of value on which it's good to have fitting attitudes. And there does seem to be something to this idea. It may seem good to admire admirable persons and objects, to intend to do things that are worth doing, to be-

²¹ Finlay (2012) comes the closest. Finlay gestures toward an interesting way to account for fit-related reasons in terms of goodness. He develops an “end-relational” semantics for “good”, according to which something is good just in case it's good for some end (where the end in question is stated explicitly or otherwise made relevant by the context). “Good for an end” is ultimately understood naturalistically—specifically, in terms of what raises the probability of the relevant end. In the case of fit-related reasons for belief, your fit-related reasons for believing that p are considerations that explain why your believing that p would be ‘epistemically good’, where epistemic goodness, on Finlay's semantics, is explained in terms of the promotion of an end—namely, the end of believing that p just in case p is true. Finlay suggests that similar stories can plausibly be told to account for fit-related reasons for other attitudes in terms of goodness. However, he only gestures at this possibility and, although interesting, this way of accounting for fit-related reasons in terms of goodness is highly committal. For one, it requires adopting Finlay's end-relational semantics for “good”.

lieve credible propositions, and to desire desirable outcomes. Of course, such attitudes needn't be of instrumental value. But such attitudes might be thought to be good for their own sakes—or of *final value*—insofar as they fit their objects.²²

So suppose fitting attitudes have final value. Then such attitudes are, presumably, *proportionally* valuable—i.e., valuable in proportion to the degree to which they're fitting.²³ Thus, given a theory of value on which fitting attitudes are of final value, facts that make an attitude to some degree *fitting*—i.e., fit-related reasons for the attitude—plausibly count as, or “add up to”, facts that make having that attitude to some degree *finally valuable*. Consequently, a theory of value according to which fitting attitudes are of final value plausibly yields the following result:

Fittingness-Value Link (FVL) Necessarily, a fact p makes ϕ -ing fitting (to some degree) only if p makes ϕ -ing finally valuable (to some degree).

FVL renders R_V extensionally consistent with the existence of fit-related reasons. But the proposal faces three serious worries. I'll raise the first two, consider a response to both, and then raise the third, to which there is no good response.

First, FVL fares poorly when it comes to cases. Recall the initial counterexamples to R_V . It's implausible that the fact that your colleague just received a promotion makes it good to envy her; or that the fact that your friend insists on talking to you throughout the movie makes it good to be annoyed with him; or that the fact that there are eight specks of dust on my desk makes it good to believe that the number of dust-specks on my desk is even.

Second, the conjunction of R_V and FVL has absurd implications. FVL entails that if a fact makes some attitude fitting to some degree, that fact makes that attitude good to some degree. And R_V entails that if a fact makes doing some act good to some degree, that fact is a reason to do that act. So together, R_V and FVL entail that whenever your doing some act would help to bring it about that you have some attitude that's fitting, that fact is a reason for you to do that act. But this is absurd. Suppose you believe now that, tomorrow, you'll have a bad day. Is the fact that it will make credible the proposition that you'll have a bad day tomorrow a reason to stab yourself in the leg today? Or suppose you fear flying. Is the fact that it will make

²² See Hurka (2001), esp. ch. 1, and Sylvan (2012).

²³ For an extended defense of this proportionality thesis, see Hurka (2001), esp. ch. 2.

your next flight genuinely fearsome a reason to tamper with the plane's engine? Together, R_V and FVL imply that the answer to both of these questions is "Yes".

Defenders of FVL might respond to each of the above worries in the following way. Mark Schroeder has argued that our intuitions to the effect that there is no reason for a response are unreliable.²⁴ When the reasons for a response are very weak, Schroeder claims, we're easily misled into thinking there is *no* reason for the response. And Schroeder's point easily extends to intuitions to the effect that there is nothing *good* about something. Perhaps when something is just a *little* good, but very bad overall, we're easily misled into thinking it's in no way good. Suppose all of this is right. Then fans of FVL can respond to each of the above worries in the following way. With regard to the first, they can accept that envying your colleague and being annoyed with your friend are bad overall, but maintain that it's also good to envy your colleague and to be annoyed with your friend—it's just that the value of having these attitudes is heavily outweighed. With regard to the second, they can accept that the reasons to stab yourself in the leg and to tamper with the plane's engine are very weak, since there are so many reasons not to do these things, but maintain that there's still *some* reason to stab yourself in the leg and to tamper with the plane—it's just that the relevant reasons are, again, heavily outweighed.

These are fair responses, and I have no quick answer to them. However I will say this. FVL seems to me a highly substantial and unattractive commitment, especially when paired with R_V . And if we can avoid making such commitments and still accommodate a view on which both fit- and value-related reasons are genuine reasons, then I think that we should. And in just a moment, I'll explain how we can do exactly this. But before I do, let me raise a third worry about the strategy currently under consideration—one which value-firsters can't so easily avoid.

As a metaphysical account of the nature of reasons, R_V isn't merely an extensional thesis—it's also an *explanatory* one. R_V entails that whenever a fact is a reason for some response, that's *because* that fact explains why it would be good to have that response. And while this seems spot on in the case of value-related reasons, this explanatory commitment is intuitively very implausible when it comes to fit-related reasons. The fact that Sharon spends a great deal of time doing charity work isn't a reason for you to admire her *because* it explains why your admiring her would be good. (And this even if, as per FVL, we think that as a matter of extensional fact, this fact *does* make your admiring Sharon good.) Indeed, the explanation for why the

²⁴ See especially Schroeder (2007: 92-7).

fact that Sharon spends much of her time doing charity work is a reason for you to admire Sharon plausibly has nothing to do with *you* and *your attitudes*. Instead, it has everything to do with *Sharon*: this fact is a reason to admire Sharon simply because it explains why she *merits* admiration, or is *worthy of* admiration, or is *fit* to admire.

So even if we grant that value-firsters who accept FVL can plausibly render R_V extensionally consistent with the existence of fit-related reasons, their approach to normativity remains explanatorily inadequate. And there is, I think, no good response to this worry. It's an essential commitment of the value-first approach that all normative facts ultimately need to be explained by appeal to—or grounded in—facts about value. But, as I've just now suggested, not all reasons-facts—i.e., facts about which facts are reasons—can be felicitously explained in this way. So at best, FVL provides value-firsters with only a *partial* response to the RKR problem: while the principle perhaps allows value-firsters to say *that* fit-related reasons are genuine reasons, it fails to secure them an adequate account of *why* such reasons are reasons.

So although this second possible response to the RKR problem is much more plausible than RKR skepticism, it fails to provide value-firsters with a fully adequate answer to the problem. And so I think neither of the ways in which value-firsters might respond to the RKR problem can plausibly succeed. The upshot: value-firsters can neither plausibly deny nor reconcile their approach to normativity with a normative view that says that fit-related reasons are genuine reasons for attitudes.

5. The Fittingness-First Approach

In what's preceded I've argued that neither the reasons- nor the value-first approach to normativity can accommodate a view on which both fit- and value-related reasons are genuine reasons. Now I'll explain how the fittingness-first approach can.

5.1 *The Fittingness-First Account of Value*

Consider first the fittingness-first account of value. According to it:

Value-in-terms-of-Fit (V_F) For p to be a respect in which x is good is for p to explain why it would be fitting (to some degree) to desire x .²⁵

²⁵ Like V_R , V_F is in the first instance an account of both final and instrumental value (and, as is the case with V_R , there are prospects for generalizing the account to goodness-for and the varieties of attributive goodness [see, e.g., McHugh and Way, 2016]). The 'to some degree'

The first thing to note about V_F is that it's not new. Accounts of value in its spirit were famously proposed by Brentano (1889/2009) and Ewing (1948), and also suggested by, e.g., Broad (1930), Brandt (1946), Wiggins (1987), and McDowell (1998).²⁶

The second thing to note about V_F is that, unlike V_R , it doesn't face the WKR problem. The fact that a demon will kill you unless you desire a cup of mud is a reason for you to desire a cup of mud—this fact explains why your desiring a cup of mud would be good. But the fact of the demon's threat doesn't also explain why your desiring a cup of mud would be *fitting*—this fact doesn't make a cup of mud fit to desire. This judgment is prevalent in the literature and certainly would have been shared by V_F 's historical defenders.²⁷ Moreover, it seems intuitively quite clear. Remember: the fittingness relation is the relation in which a response stands to an object when the object merits—or is worthy of—that response. And while there's certainly *something* to be said for desiring a cup of mud in circumstances in which you'll be killed if you don't, it would seem infelicitous to say that the mud *merits* being desired, or that it's *worthy of* being desired. Rather, it seems, the fitting response to have toward such an object would be closer to indifference, or perhaps disgust.

Still, it might be suggested, perhaps there is a 'wrong kind of fittingness', or a notion of 'pragmatic fittingness', such that, in this sense, it *would* be fitting to desire the mud. If so, then V_F would face a problem parallel to the WKR problem: the demon's threat would make it (pragmatically) fitting to desire a cup of mud, although

qualifier is necessary to accommodate the fact that, in many cases, a respect in which something is good will make wanting that thing fitting *to some degree*, though not fitting *overall*.

²⁶ Fittingness-based accounts of value are also defended by more contemporary authors, including Danielsson and Olson (2007), Chappell (2012), and McHugh and Way (2016). As I indicated earlier, Chappell and McHugh and Way defend fittingness-first views similar in spirit to that proposed in this paper (though Chappell is specifically concerned to defend the conceptual thesis that 'fittingness' is the sole primitive normative concept, whereas McHugh and Way and myself are concerned to defend the metaphysical thesis that *fittingness* is the fundamental normative relation). There are a number of interesting and consequential differences among contemporary fittingness-firsters, but perhaps the most crucial lies in our importantly dissimilar accounts of *reasons*. More on this in the section that follows.

²⁷ Here's a non-exhaustive list of authors who endorse this judgment, either explicitly or tacitly: Brandt (1946); Brentano (1889/2009); Broad (1930); Chappell (2012); Chisholm (1986); D'Arms and Jacobson (2000b); Danielsson and Olson (2007); Ewing (1948); McDowell (1998); McHugh and Way (2016); Rabinowicz and Rønnow-Rasmussen (2004); Sharadin (2015); Schroeder (2010); Svavarsdottir (2014); Thomson (2008); Wiggins (1987); Zimmerman (2011).

the demon’s threat wouldn’t also make a cup of mud valuable or good. Counterexample. Call this the “Wrong Kind of Fittingness” problem (WKF problem).

The “WKF problem” is not, I think, a problem. The notion of ‘pragmatic fittingness’ invoked in its statement seems to me spurious.²⁸ To say that it would be pragmatically fitting to want x is, I think, just to say that it would be instrumentally good to want x —or, on my account, that it would be *fitting to want* to want x . But, as should by now be clear, facts that make it instrumentally good to want x needn’t also make it fitting to want x , any more than facts that make it instrumentally good to admire a deplorable dictator need also make it fitting to admire the dictator.

So V_F doesn’t face the WKR problem (or the WKF problem). But by itself, V_F doesn’t provide the resources necessary to accommodate and explain the existence of both fit- and value-related reasons. This is because, by itself, V_F implies nothing about reasons. What we need, then, is a fittingness-based account of reasons.

5.2 The Fittingness-First Account of Reasons

Three fittingness-based accounts of reasons are already suggested in the literature:

- (1) For p to be a reason to ϕ is for p to be evidence that it is fitting to ϕ (Thomson 2008).
- (2) For p to be a reason to ϕ is for p to explain why it is fitting to ϕ (Chappell 2012).
- (3) For p to be a reason to ϕ is for p to be a premise of a fittingness-preserving pattern of reasoning from fitting responses to ϕ -ing (McHugh and Way 2016).

The main problem with each of (1)-(3) is the same: all three accounts are incompatible with the existence of value-related reasons for attitudes. (1) and (2) entail that something is a reason for a response only if it explains why (or is evidence that) it would be fitting to have that response. But, as we’ve already seen, not all value-related reasons for attitudes also explain why (or are evidence that) it would be fitting to have those attitudes. (3) entails that something is a reason for a response only if it’s a premise of a fittingness-preserving pattern of reasoning from fitting responses to that response. But not all value-related reasons will satisfy this condition. For example, reasoning from the belief that the demon will kill you unless you desire a

²⁸ What could it mean to say that something is ‘pragmatically worthy of’ a certain response?

cup of mud to desiring a cup of mud doesn't preserve fittingness, since desiring a cup of mud isn't fitting. So, according to (3), the fact that the demon will kill you unless you desire a cup of mud isn't a reason for you to desire a cup of mud. So none of (1)-(3) are compatible with the existence of value-related reasons for attitudes.

So how might an adequate fittingness-based account of reasons go? Recall that I began this paper with the observation that reasons for attitudes are facts that count in favor of those attitudes, but a fact can count in favor of an attitude in two different ways: by making the attitude fitting or by making the attitude good to have. This observation suggests that reasons for attitudes come in two different kinds: fit-and value-related. The accounts of reasons so far considered in this paper have been able to account for one or the other of these kinds of reasons, but none has been able to account for both. The value-firster's account of reasons easily accounts for value-related reasons, but founders when it comes to fit-related reasons. And each of the above fittingness-based accounts of reasons accounts for the existence of fit-related reasons, but fails to account for value-related reasons. So what we need is an account of reasons that accounts for both fit- *and* value-related reasons, an account which predicts that—and explains why—a fact gets to be a reason *either* when it makes an attitude fitting *or* when it makes an attitude good. The fittingness-firster has the resources necessary to provide such an account. So consider:

Reasons-in-terms-of-Fit (R_F) For p to be a reason to φ is for p to explain either why it is fitting to φ or why it is fitting to want to φ .

By itself, R_F entails that whenever a fact makes some attitude fitting, that fact is a reason for that attitude. So the fact that fit-related reasons are genuine reasons for attitudes follows easily. But paired with V_F , R_F also entails that value-related reasons are genuine reasons. This is easy to see. V_F entails that whenever a fact makes some attitude good to have, that fact makes that attitude fitting to want. And R_F entails that whenever a fact makes some attitude fitting to want, that fact is a reason for that attitude. So V_F and R_F jointly entail that value-related reasons are genuine reasons for attitudes. And so together with V_F , R_F predicts that—and explains why—both fit- and value-related reasons are genuine reasons.

Besides this, R_F has two further, notable attractions. First, recall that in section 3, I highlighted the attractiveness of the thesis that fit-related reasons to want an attitude count as, or “add up to”, reasons for that attitude—i.e., that p is a fit-related reason to want some attitude only if p is a reason for that attitude. I motivat-

ed this thesis by appealing to the plausibility of its parallel in the case of reasons to act—i.e., that p is a fit-related reason to want to do some act only if p is a reason to do that act. And I noted that, if true, these theses shouldn't be taken as brute truths—that, if true, we ultimately need to be able to provide some explanation of *why* they're true (see note 13). But now, notice that R_F allows us to do this: facts that make it fitting to want some response (whether it be an act or attitude) are reasons for that response, because it lies in the nature of reasons that if a fact makes some response fitting to want, then that fact is a reason for that response.

Next, recall that in section 4, I charged the value-first approach with explanatory inadequacy. I argued that while the value-firster's commitment to grounding all reasons-facts in facts about value looks plausible when it comes to value-related reasons, this commitment yields intuitively implausible results when it comes to fit-related reasons. Thus, although it's plausible that the explanation for why the fact that the dictator will kill you unless you admire him is a reason for you to admire the dictator has to do with the value of your admiring him, it's not at all plausible that the explanation for why the fact that Sharon spends much of her time doing charity work is a reason for you to admire Sharon has to do with the value of your admiring her. But now notice R_F doesn't face this difficulty and, in fact, that it gets the best of both worlds. For the fittingness-firster who accepts R_F can say that what makes the fact that the dictator will kill you unless you admire him a reason for you to admire the dictator is that this fact explains why it would be fitting to want—i.e., *good*—to admire him. And she can say that what makes the fact that Sharon spends much of her time doing charity work a reason for you to admire Sharon is that this fact explains why she *merits* admiration, or is *worthy of* admiration, or is *fitting* to admire.²⁹

6. Conclusion

In the first part of this paper, I argued that neither the reasons- nor the value-first approach to normativity can accommodate a normative view on which both fit- and value-related reasons are genuine reasons. In the second part of the paper, I explained how my version of the fittingness-first approach can accommodate such a

²⁹ R_F has further attractions as well, which, although shared by the other accounts of reasons considered in this paper, are worth noting here. For example, R_F offers a *unified* account of reasons, in the sense that it applies to reasons for responses of all types, i.e., to reasons to act, believe, desire, and so forth. And R_F easily accounts for the existence of *outweighed* reasons, since a response can be fitting to some degree, without being fitting overall, or most fitting.

view, and argued that alternative, existing fittingness-first views cannot. Because my version of the fittingness-first approach similarly accounts for all other normative properties and relations in terms of just a single normative relation, it's no less parsimonious than the reasons-first approach, the value-first approach, or existing versions of the fittingness-first approach. And since, unlike any of these alternatives, my version of the fittingness-first approach can plausibly accommodate and explain the existence of all of the reasons there are, it provides a substantively superior and so overall more attractive ontology of normativity. So, as compared to the reasons- and value-first approaches, as well as existing versions of the fittingness-first approach, my version of the fittingness-first approach ought to be preferred.

Of course, much more would need to be done for a full defense of the fittingness-first view I've advanced in this paper. I've argued only that the view should be preferred to prominent alternatives, not that it should be accepted full stop. A more thorough defense is a task I leave for future work. For now, I hope only to have offered a strong initial case for my fittingness-first view, and thus to have made the picture of normativity that it provides seem worthy of further investigation.³⁰

References

- Alston, W. *Beyond Justification: Dimensions of Epistemic Evaluation*. Ithaca: Cornell University Press, 2005.
- Brandt, R. "Moral Valuation." *Ethics* 56 (1946): 106-21.
- Brentano, F. *The Origin of Our Knowledge of Right and Wrong*. Oxford: Routledge, 1889/2009.
- Broad, C.D. *Five Types of Ethical Theory*. London: Routledge, 1930.
- Broome, J. *Rationality Through Reasoning*. Hoboken, NJ: Wiley-Blackwell, 2013.
- . "Reasons." In *Reason and Value: Themes from the Moral Philosophy of Joseph Raz*, edited by J. Wallace, M. Smith, S. Scheffler, and P. Pettit, 228-55. Oxford: Oxford University Press, 2004.
- . "Reply to Southwood, Kearns and Star, and Cullity." *Ethics* 119 (2008): 96-108.
- Brunero, J. "Reasons, Evidence, and Explanations." In *The Oxford Handbook of Reasons and Normativity*, edited by D. Star. Oxford: Oxford University Press, forthcoming.

³⁰ Thanks to (many, many people).

- Chappell, R. "Fittingness: The Sole Normative Primitive." *Philosophical Quarterly* 62 (2012): 684-704.
- Chisholm, R. *Brentano and Intrinsic Value*. Cambridge: Cambridge University Press, 1986.
- Crisp, R. "Review of John Kupperman, *Value...and What Follows*." *Philosophy* 75 (2000): 458-62.
- D'Arms, J. and D. Jacobson. "The Moralistic Fallacy: On the 'Appropriateness' of Emotions." *Philosophy and Phenomenological Research* 61 (2000a): 65-90.
- . "Sentiment and Value." *Ethics* 110 (2000b): 722-48.
- Danielsson, S. and J. Olson. "Brentano and the Buck-Passers." *Mind* 116 (2007): 511-22.
- Ewing, A.C. *The Definition of Good*. London: Routledge & Kegan Paul, 1948.
- Feinberg, J. *Doing & Deserving: Essays in the Theory of Responsibility*. Princeton, NJ: Princeton University Press, 1970.
- Fine, K. "Essence and Modality." *Philosophical Perspectives* 8 (1994): 1-16.
- Finlay, S. "Explaining Reasons." *Deutsche Jahrbuch für Philosophie* 4 (2012): 112-26.
- . "The Reasons that Matter." *Australasian Journal of Philosophy* 84 (2006): 1-20.
- Gibbard, A. *Wise Choices, Apt Feelings*. Cambridge, MA: Harvard University Press, 1990.
- Goldman, A. *Knowledge in a Social World*. Oxford: Oxford University Press, 1999.
- Howard, C. "In Defense of the Wrong Kind of Reason." *Thought: A Journal of Philosophy* 5 (2016a): 53-62.
- . "Transparency and the Ethics of Belief." *Philosophical Studies* 173 (2016): 1191-1201.
- Hurka, T. *Virtue, Vice, and Value*. Oxford: Oxford University Press, 2001.
- Hieronymi, P. "The Wrong Kind of Reason." *Journal of Philosophy* 102 (2005): 437-57.
- Jackson, F. *From Metaphysics to Ethics*. Oxford: Oxford University Press, 1997.
- Kavka, G. "The Toxin Puzzle." *Analysis* 43 (1983): 33-36.
- Kearns, S. and D. Star. "Reasons: Explanations or Evidence?" *Ethics* 119 (2008): 31-56.
- . "Reasons as Evidence" In *Oxford Studies in Metaethics*, vol. 4, edited by R. Shafer-Landau, 215-42. Oxford: Oxford University Press, 2009
- Lynch, M. *True to Life: Why Truth Matters*. Cambridge, Massachusetts: MIT Press, 2004.
- Maguire, B. "There are no Reasons for Emotions or Desires." Unpublished manuscript, The University of North Carolina at Chapel Hill, 2016.

- Markovits, J. "Internal Reasons and the Motivating Intuition." In *New Waves in Metaethics*, edited by M. Brady, 141-165. Basingstoke: Palgrave Macmillan, 2011.
- McDowell, J. "Values and Secondary Qualities." Reprinted in his *Mind, Value and Reality*. Cambridge, MA: Harvard University Press, 1998.
- McHugh, C. and J. Way. "Fittingness First." *Ethics* 126 (2016): 575-606.
- McKeever, S. and M. Ridge. "Elusive Reasons." In *Oxford Studies in Metaethics*, vol. 7, edited by R. Shafer-Landau, 110-37. Oxford: Oxford University Press, 2012.
- Moore, G.E. *Principia Ethica*. Cambridge: Cambridge University Press, 1903.
- Papineau, D. "There are No Norms of Belief." In *The Aim of Belief*, edited by T. Chan, 64-79. Oxford: Oxford University Press, 2013.
- Parfit, D. *On What Matters*. Oxford: Oxford University Press, 2011.
- Piller, C. "Content-Related and Attitude-Related Reasons for Preferences." *Royal Institute of Philosophy Supplement* 81 (2006): 155-81.
- 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. *Engaging Reason*. Oxford: Oxford University Press, 1999.
- . *From Normativity to Responsibility*. Oxford: Oxford University Press, 2011.
- . "Reasons: Practical and Adaptive." In *Reasons for Action*, edited by D. Sobel and S. Wall, 37-57. Cambridge: Cambridge University Press, 2009.
- Reisner, A. "The Possibility of Pragmatic Reasons for Belief and the Wrong Kind of Reason Problem." *Philosophical Studies* 145 (2009): 257-72.
- . "Weighing Pragmatic and Evidential Reasons for Belief." *Philosophical Studies* 138 (2008): 17-27.
- Rosen, G. "The Alethic Conception of Moral Responsibility." In *The Nature of Moral Responsibility: New Essays*, edited by R. Clarke, M. McKenna, and A.M. Smith, 65-88. Oxford: Oxford University Press, 2015a.
- . "Real Definition." *Analytic Philosophy* 56 (2015b): 189-209.
- Ross, W.D. *Foundations of Ethics*. Oxford: Clarendon Press, 1939.
- . *The Right and the Good*. Oxford: Clarendon Press, 1930.
- Scanlon, T. *What We Owe to Each Other*. Cambridge: Harvard University Press, 1998.
- Schroeder, M. *Slaves of the Passions*. Oxford: Oxford University Press, 2007.

- . "Value and the Right Kind of Reason." In *Oxford Studies in Metaethics*, vol. 5, edited by R. Shafer-Landau, 25-55. Oxford: Oxford University Press, 2010.
- . "The Ubiquity of State-Given Reasons." *Ethics* 122 (2012): 457-88.
- Shah, N. "A New Argument for Evidentialism." *Philosophical Quarterly* 56 (2006): 481-98.
- Sharadin, N. "Reasons Wrong and Right." *Pacific Philosophical Quarterly* 96 (2015): 1-29.
- . "Schroeder on the Wrong Kind of Reasons Problem for Attitudes." *Journal of Ethics and Social Philosophy* 7 (2013): 1-8.
- Skorupski, J. *The Domain of Reasons*. Oxford: Oxford University Press, 2010.
- Smith, M. "Two Kinds of Consequentialism." *Philosophical Issues* 19 (2009): 257-72.
- Snedegar, J. "Negative Reason Existentials." *Thought* 2 (2013): 108-116.
- Steglich-Peterson, A. "How to be a Teleologist about Epistemic Reasons." In *Reasons for Belief*, edited by A. Steglich-Peterson and A. Reisner, 13-33. Cambridge: Cambridge University Press, 2011.
- Stratton-Lake, P. "How to Deal with Evil Demons: Comment on Rabinowicz and Rønnow-Rasmussen." *Ethics* 115 (2005): 788-98.
- Svavarsdottir, S. "Having Value and Being Worth Valuing." *Journal of Philosophy* 111 (2014): 84-109.
- Sylvan, K. "Truth Monism Without Teleology." *Thought: A Journal of Philosophy* 1 (2012): 161-63.
- Thomson, J.J. *Normativity*. Chicago and La Salle: Open Court, 2008.
- Way, J. "Transmission and the Wrong Kind of Reason." *Ethics* 122 (2012): 489-515.
- . "Value and Reasons to Favour." In *Oxford Studies in Metaethics*, vol. 8, edited by R. Shafer-Landau, 27-49. Oxford: Oxford University Press, 2013.
- Wedgwood, R. *The Nature of Normativity*. Oxford: Oxford University Press, 2007.
- Wiggins, D. "A Sensible Subjectivism?" In his *Needs, Values, Truth: Essays in the Philosophy of Value*. Oxford: Blackwell, 1987.
- Williams, B. "Internal and External Reasons." Reprinted in his *Moral Luck*. Cambridge: Cambridge University Press, 1981.
- Zimmerman, M. "Partiality and Intrinsic Value." *Mind* 120 (2011): 447-83.
- . "Varieties of Moral Responsibility." In *The Nature of Moral Responsibility: New Essays*, edited by R. Clarke, M. McKenna, and A.M. Smith, 45-65. Oxford: Oxford University Press, 2015.