THE FUNDAMENTALITY OF FIT

by

Christopher M. Howard

A Dissertation Submitted to the Faculty of the

DEPARTMENT OF PHILOSOPHY

In Partial Fulfillment of the Requirements

For the Degree of

DOCTOR OF PHILOSOPHY

In the Graduate College

THE UNIVERSITY OF ARIZONA

2017

THE UNIVERSITY OF ARIZONA
GRADUATE COLLEGE

As members of the Dissertation Committee, we certify that we have read the dissertation prepared by Christopher Howard, titled *The Fundamentality of Fit* and recommend that it be accepted as fulfilling the dissertation requirement for the Degree of Doctor of Philosophy.

____ Date: July 27, 2017
Mark Timmons

____ Date: July 27, 2017
Stewart Cohen

____ Date: July 27, 2017
Terence Horgan

____ Date: July 27, 2017
Connie Rosati

Final approval and acceptance of this dissertation is contingent upon the candidate's submission of the final copies of the dissertation to the Graduate College.

I hereby certify that I have read this dissertation prepared under my direction and recommend that it be accepted as fulfilling the dissertation requirement.

____ Date: July 27, 2017
Dissertation Director: Mark Timmons

STATEMENT BY AUTHOR

This dissertation has been submitted in partial fulfillment of the requirements for an advanced degree at the University of Arizona and is deposited in the University Library to be made available to borrowers under rules of the Library.

Brief quotations from this dissertation are allowable without special permission, provided that an accurate acknowledgement of the source is made. Requests for permission for extended quotation from or reproduction of this manuscript in whole or in part may be granted by the copyright holder.

SIGNED: Christopher M. Howard

ACKNOWLEDGMENTS

This section would be far too long were I to list everyone to whom I owe a debt of gratitude for assisting (in various ways) in the production of the work that constitutes this dissertation. Still, I would like to acknowledge the following people:

First, my dissertation committee: Mark Timmons, Stewart Cohen, Terence Horgan, and Connie Rosati. Thank you for your continuous and enormously helpful guidance over the course of my time in graduate school.

Second, my peers at the University of Arizona, and in particular Sameer Bajaj, Jacob Barrett, Victor Kumar, Santiago de Jesus Sanchez Borboa, and Eyal Tal. I'll never be able to repay you for the countless hours of discussion, all of which significantly benefited this work.

Third, Jonathan Way and Selim Berker. Thank you for taking a bet on me early in my career, and for all the time you spent helping me to become a better philosopher.

And finally, my loving partner, Kristen Howard, whose personal and practical support has helped this dissertation reach its full potential.

DEDICATION

To my parents, Andrew and Sharon Howard, whose relentless and selfless support I have continually relied on in the course of creating this work. I simply can’t thank you enough.

TABLE OF CONTENTS

ABSTRACT

Normative reasons for attitudes are facts that count in favor of those attitudes, but a fact can favor an attitude in two very different ways. One way in which a fact can favor an attitude is by making the attitude *fitting* (apt, merited, correct). For example, the fact that Sharon spends much of her time doing charity work is a fact that favors admiring Sharon, since it's a fact that makes her admirable, and so fit to admire. Call any fact that favors an attitude by making it fitting a "fit-related reason."

A second way in which a fact can favor an attitude is by making the attitude somehow valuable, or good to have. For example, the fact that an evil dictator will order my execution unless I admire him is a fact that favors my admiring the dictator, since it's a fact that makes my admiring him good. Call any fact that favors an attitude by making it somehow good to have a "value-related reason."

This dissertation has two main goals. The first is to develop an ontology of normativity that can accommodate a view on which both fit- and value-related reasons are genuine reasons. Many authors, including Derek Parfit, T.M. Scanlon, and Mark Schroeder, favor a "reasons-first" ontology of normativity: they take reasons to be normatively basic, and claim that all other normative facts, properties, and relations can be analyzed or accounted for in terms of the reason relation. A central alternative, famously defended by G.E. Moore in *Principia Ethica*, is the "value-first" ontology—an ontology that takes value or goodness to be normatively basic and claims that the rest of the normative can be accounted for in terms of the property of being good. In the opening chapter of my dissertation, "The Fundamentality of Fit," I advance an ontology of normativity, originally suggested by Franz Brentano and A.C. Ewing, according to

which *fittingness* is the basic normative relation, in terms of which the rest of the normative can be explained. I argue that neither the reasons- nor the value-first ontology can accommodate a view on which both fit- and value-related reasons are genuine reasons. Then I explain how my "fittingness-first" ontology can.

Of course, any threat to the plausibility of the view that both fit- and value-related reasons are genuine reasons would undermine the case for my ontology of normativity. And so a full defense of my fittingness-first ontology will require a systematic defense of the substantive normative view it's designed to accommodate. The second goal of my dissertation is to provide this defense.

A normative view that says that both fit- and value-related reasons are genuine reasons consists in three component claims: (1) that fit-related reasons are genuine reasons; (2) that value-related reasons are genuine reasons; and (3) that fit- and value-related reasons can be compared against one another to yield univocal verdicts concerning what attitudes one ought, all-things-considered, to have. The first of these claims—that fit-related reasons are genuine reasons—is among the most widely shared in contemporary normative theory. The latter two, however, are more controversial. In the second and third chapters of this dissertation, I defend each of these claims in turn.

One way to showcase the plausibility of a normative view that says that both fit- and value-related reasons are genuine reasons is to show that it explains our intuitions in a variety of substantive normative debates. This would, in turn, provide support for my fittingness-first ontology, since, relative to its main competitors, my ontology uniquely accommodates and predicts such a view. In the final chapter, I put this methodological observation into practice by testing the substantive normative predictions of my fittingness-first view against our intuitions in the debate concerning what kinds of considerations can provide reasons for love. I argue that acknowledging the existence of both fit- and value-related reasons for love solves a number of persistent problems in this debate.

CHAPTER 1: The Fundamentality of Fit[1]

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

---

[1] This essay won the 2017 Marc Sanders Prize in Metaethics and is forthcoming in *Oxford Studies in Metaethics*, vol. 14, ed. Russ Shafer-Landau (Oxford University Press, forthcoming).

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.[2] 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

[2] 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).

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.[3] 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.[4]

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 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.[5]

---

[3] For the rest of this paper, unless I indicate otherwise, "reason" means "normative reason."

[4] 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.

[5] 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.

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*.[6]

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

[6] 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.

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.[7]

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*.[8]

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

[7] 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 reasons 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.

[8] 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.

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.[9] 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.

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.[10]

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

---

[9] This example famously comes from Crisp (2000).
[10] See D'Arms and Jacobson (2000a) and Rabinowicz and Rønnow-Rasmussen (2004).

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*.[11] 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 (D'Arms and Jacobson 2000b, 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 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"

[11] I borrow this characterization of "incentives" from Way (2012, 492.)

(Danielsson and Olson, 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.[12] 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.[13] 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 (Danielsson and Olson 2007, Markovits 2011, Rabinowicz and Rønnow-Rasmussen 2004).[14] 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

[12] 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.

[13] 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).

[14] 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.

2013). But this line of argument is unavailable to reasons-firsters, since these theorists are committed to there being reasons that *violate* a response-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.[15] 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.[16] And so defenders of V_R are committed to there being reasons that violate a response-condition.[17] 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.[18]

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

[15] 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."

[16] 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).

[17] Way (2012, 512*f*) 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 (chapter 2 of this dissertation) 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.

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

offer reasons-firsters an adequate answer to the WKR problem. And so WKR skepticism is, I think, untenable as a response to this problem.

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 the 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 solution 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.[19]

[19] 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

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.

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.[20]

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

---

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).

[20] 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.

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.[21] 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 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.

[21] 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.

Now, as far as I know, no value-firster has offered a systematic response to the RKR problem in print.[22] 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 *enviable*, 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

[22] 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."

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 believe 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.[23]

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.[24] 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

[23] See Hurka (2001), esp. ch. 1, and Sylvan (2012).
[24] For an extended defense of this proportionality thesis, see Hurka (2001), esp. ch. 2.

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 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.[25] 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

[25] See especially Schroeder (2007, 92-7).

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 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*.[26]

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).[27]

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

[26] 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" 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*.

[27] 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.

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.[28] 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 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.[29] 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

[28] 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).

[29] What could it mean to say that something is "pragmatically worthy of" a certain response?

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 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 motivated 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.[30]

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 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

[30] 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.

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.

CHAPTER 2: In Defense of the Wrong Kind of Reason[1]

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 (RKRs): considerations that count in favor of ϕ-ing which *do* bear on whether ϕ-ing would be correct. Examples of RKRs are

[1] This paper appears in *Thought: A Journal of Philosophy* 5 (2016): 53-62.

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.[2]

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

[2] 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.

you depends on whether you *actually* come to desire a cup of mud. Then, as some have noted,[3] 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. 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 WKRs transmit in a way that's distinctive of such reasons—i.e., in a way that's different from RKRs. Way suggests that this fact calls for explanation, that an account of WKRs 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 WKRs transmit is *not* distinctive—that RKRs 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 WKRs transmit: the way in which such reasons transmit isn't a feature of WKRs in particular but rather of reasons for attitudes *in general*. Thus WKR defenders can explain the way in which WKRs transmit by appeal to

[3] See Danielsson and Olson (2007), 513-14; and for similar concerns about the lack of argument for WKR skepticism, see Rabinowicz and Rønnow-Rassmussen (2004).

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 WKRs 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, 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.[4]

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

To bear this out, consider first the case of RKRs 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 RKRs 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 RKRs to desire transmit in this way, or for why RKRs for other attitudes transmit in a parallel fashion (i.e. across

[4] 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), 494-96.

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 WKR skeptics can unproblematically appeal in arguing for WKR skepticism, without first needing to explain.[5]

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.[6]

[5] 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 July 3, 2017). In their comment, Rabinowicz and Rønnow-Rasmussen propose that we might explain the way in which RKRs transmit by appeal to facts about the transmission of value among objects of attitudes for which there are RKRs. This proposal looks promising (at least when it comes to RKRs 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.

[6] This example comes from Way (2012), 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

Now remember: according to Way, the fact that WKRs 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 RKRs, 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.[7] 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 WKRs 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.

---

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.

[7] 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], 500-2).

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

> **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, 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, 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.[8] 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

[8] Cf. Hurka (2001), esp. ch. 1, and Sylvan (2012).

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.[9] 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.[10] 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*.[11] And these look like compelling axiological results. But what's

[9] 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.

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

[11] 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 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.

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 which WKRs 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 RKRs might differ in two ways: WKRs might transmit in a way that RKRs don't or WKRs might fail to transmit in a way that RKRs do. I've argued that the transmission profiles of WKRs and RKRs don't differ in the first way: that RKRs transmit in the way that WKRs do, viz., across facilitative connections between attitudes. However, I do think they differ in the second way: that RKRs transmit in a way that WKRs don't, viz., across facilitative connections between the objects of attitudes.[12] 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*.

[12] Way (2012, 497) notes this difference in the transmission profiles of WKRs and RKRs in passing, but doesn't make any hay out of it.

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.[13] So the Wrong Desire Pattern doesn't hold.

To generalize: whereas RKRs 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 RKRs 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, albeit one which takes as a premise a different difference in the ways in which WKRs and RKRs transmit.

And in fact, WKR skeptics can plausibly explain the difference in question—they can plausibly explain why, unlike RKRs, 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 RKRs to want attitudes—i.e., RKRs for higher-order desires—fail to transmit. To see this, consider the

[13] I borrow this example from Way (2012, 497).

> **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 RKRs 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 RKRs 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 RKRs for higher-order desires, and RKRs for higher-order desires fail to transmit according to such patterns.

So WKR skeptics can plausibly explain why, unlike RKRs, 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 RKRs 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 RKRs 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 RKRs for higher-order desires, and RKRs 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 RKRs 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.[14] Why, then, shouldn't we think that a parallel thesis holds in the case of reasons for attitudes? If RKRs to want to do some act are reasons to do that act, then why shouldn't RKRs to want some attitude be reasons for that attitude? True, WKR defenders who accept this thesis ultimately owe an explanation of its truth.[15] 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 RKRs to want attitudes

[14] 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).

[15] On this, see Way (2012). I offer such an explanation in Howard (forthcoming; see chapter 1 of this dissertation). But for reasons of space, I can't discuss the details here.

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.

CHAPTER 3: Weighing Epistemic and Practical Reasons for Belief

1. Introduction

Considering the contemporary debate concerning what kinds of considerations can provide normative reasons for belief, we find two extremes. At one extreme lies *Alethism*, according to which the only kinds of considerations that can provide reasons for (or against) belief are *epistemic*: considerations that are in some way relevant to getting at the truth and avoiding error (Shah 2006, Thomson 2008, Parfit 2011, Whiting 2014, Way 2016).[1] At the other extreme lies *Robust Pragmatism*, according to which only practical, broadly benefit-related, considerations can provide genuine reasons for believing (Papineau 2013, Rinard 2015).[2]

I believe that both these views are wrong and that, as is often the case with extremes, the truth lies in between. Normative reasons for a belief are considerations that are relevant to determining whether one ought to have that belief, all-things-considered. And both epistemic *and* practical considerations seem entirely relevant to determining what beliefs one ought, all-things-considered, to have. It's an attractive normative

[1] I borrow the name "Alethism" and this way of characterizing the view from Leary (forthcoming). Many Alethists are *Evidentialists*—they hold that all reasons for belief are epistemic, and also that all epistemic reasons are evidence-based. But Alethism doesn't require this commitment.

[2] I borrow the name 'Robust Pragmatism' from Rinard (2015). Rinard counts moral considerations among practical reasons for belief, since the benefit indicated or (partly) grounded by a practical consideration might be a benefit to another, as opposed to oneself. I follow Rinard in this here.

hypothesis that one ought normally to have only those beliefs that are adequately supported by one's epistemic reasons. But in situations in which one's having an epistemically unsupported belief is necessary to bring about some great good (or to prevent some great evil) it seems that, all-things-considered, one ought to have that belief, despite its being epistemically defective.

That isn't to say that a normative view on which both epistemic and practical considerations can provide reasons for belief is, or even should be, incontestable. The current popularity of the opposing positions mentioned above clearly suggests otherwise. Still, this more moderate "third way" in the debate concerning what kinds of considerations can provide reasons for belief no doubt deserves a name. So let's call it *Inclusivism*. To state the view succinctly: Inclusivism claims that both epistemic and practical considerations can provide genuine reasons for belief, and that both types of consideration can contribute to (metaphysically) determining what beliefs one ought, all-things-considered, to have.[3]

This paper answers a challenge for Inclusivism. The challenge is to provide a plausible model for comparing epistemic and practical reasons for belief against one another when the two types of reasons conflict. Some suggest this can't be done (Steglich-Peterson 2011, Chappell 2012, Papineau 2013, Berker MS). They suspect that there can be no plausible model for determining what one ought to believe *all-things-*

[3] With the exception of Pascal (1670) and James (1897), and perhaps Foley (1993, 2003), this view is rarely, if ever, explicitly endorsed in print. On one reading, Marušić (2016) can be interpreted as accepting the view, though he denies this in conversation. Reisner (2008) goes some way toward defending Inclusivism (what he calls "The Mixed View") but never officially endorses it. Markovits (2010, 2014) also makes some apparently sympathetic remarks, but, again, never makes it official.

considered when there are conflicting epistemic and practical reasons.[4] I disagree. The present paper provides such a model.

The contemporary ethics of belief literature doesn't provide much precedent when it comes to models for comparing epistemic and practical reasons. The only exception of which I'm aware is a model suggested in a paper by Andrew Reisner (2008). But I'll argue that my model is preferable to Reisner's—that it has all the latter's advantages, but none of its costs.[5]

Let me be clear about my aims. First, I have no ambition to offer a full defense of Inclusivism in this paper. For one, that would require defending the still-minority view that practical considerations can provide reasons for belief. I've undertaken this task elsewhere, but won't continue the fight here (Howard 2016a, 2016b). In what follows I will simply be assuming, without argument, that practical reasons for belief exist. Nonetheless, this paper constitutes a contribution to the debate concerning what kinds of considerations can provide reasons for belief, inasmuch as it succeeds in partially clearing the way for Inclusivism—in particular, by answering what's perhaps *the* central challenge for the view.

Second, I won't discuss in detail the possible (indeed, actual) view that both epistemic and practical considerations can provide reasons for belief, but that these two types of consideration don't contribute to determining a single, all-things-considered ought and instead determine, respectively, an *epistemic ought* and a *practical ought*,

---

[4] All but one of these authors provide no argument that epistemic and practical reasons can't be compared against each other—they merely report the intuitive difficulty of seeing how such comparisons could be possible (one tends to encounter this in informal discussion of the matter, too). Berker (MS) is the exception. He surveys some existing and possible models and argues that none are satisfactory. Berker's work will feature prominently in what follows.

[5] In the historical literature, a possible model for comparing epistemic and practical reasons for belief is suggested by some remarks made by William James (1897). However, as I explain below, there are good (epistemic) reasons to think that James's model is unsatisfactory (see note 15).

which can't be compared (Feldman 2000).[6] I omit discussion of this view for the following reason. There are, as far as I can see, just two possible rationales for such a view. The first is global skepticism about the notion of an all-things-considered ought (Copp 1997, Tiffany 2007, Baker MS). Philosophers who are globally skeptical about the notion of an all-things-considered ought don't just deny that there could be an all-things-considered ought that takes into account both epistemic and practical reasons; instead, they take the notion of an all-things-considered ought to be incoherent across the board. So such authors deny, for example, that it can make sense to ask what one ought to do, all-things-considered, when there are conflicting moral and prudential reasons. I don't aim to address this sort of skepticism here. I assume that the notion of an all-things-considered ought is coherent.[7]

The second possible rationale for the kind of view under consideration is a local skepticism about the possibility of an all-things-considered ought which takes into account both epistemic and practical reasons. The source of this sort of skepticism is, presumably, doubt about the possibility of making principled and meaningful comparisons among the two types of reasons. But this brand of skepticism is precisely the target of the current paper.

Here's a roadmap. In the next section I explain why epistemic and practical reasons for belief can't plausibly be compared by a simple weighing procedure. The discussion reveals two desiderata that any model for comparing epistemic and practical reasons should satisfy. In section 3, I explain the virtues of Reisner's model for

---

[6] Strictly speaking, I think it's a (fairly common) mistake to talk about comparing *oughts*, as opposed to comparing the *reasons* underlying each respective ought. Here, I've stated the position under consideration in the former way, since this is how those who endorse it tend to state it.

[7] For defense of the all-things-considered ought, see esp. McLeod (2001); Thomson (2008).

comparing the two types of reasons, and also why his model ultimately fails. I then present my model for comparing epistemic and practical reasons (in section 4). In section 5, I conclude by considering an objection.

2. Why Simple Weighing Won't Work

Epistemic and practical reasons can't plausibly be compared by a simple weighing procedure. We can't simply combine the epistemic and practical reasons for and against a given doxastic alternative[8] in order to determine the total reason (*simpliciter*) in favor of that alternative, and then compare the total reason in favor of each alternative in order to determine which doxastic option there is decisive (or sufficient) all-things-considered reason to take.[9] Such a procedure may seem a natural choice for making comparisons among other types of reasons—for example, moral and prudential reasons for action. But for the purpose of comparing epistemic and practical reasons for belief, a simple weighing procedure simply doesn't suffice. There are two principal reasons for this, both of which are instructive.

[8] I'll assume throughout that there are three alternative positive doxastic attitudes that one might take toward a given proposition: belief, disbelief, and suspension of judgment. I realize that it's still somewhat (though not very) controversial that suspension of judgment constitutes a positive doxastic attitude. But for a powerful and rigorous defense of this claim, see Friedman (2013).

[9] Exactly how to combine the reasons at the first step of this procedure is a difficult and controversial question; for some insightful discussion, see Nair (2016). The question of how to compare the total reason in favor of each alternative at the second step is also controversial. One might opt to make such comparisons in a maximizing way, such that one has decisive all-things-considered reason for a given doxastic attitude just in case there's more total reason *simpliciter* in favor of the attitude than there is in favor of any of its alternatives, has sufficient all-things-considered reason for the attitude just in case there's at least as much total reason *simpliciter* in favor of the attitude as there is in favor of any of its alternatives, and otherwise lacks sufficient all-things-considered reason for the attitude. But one might also make the relevant comparisons in a satisficing way such that one has sufficient all-things-considered reason for a given doxastic attitude just in case the total reason *simpliciter* in favor of the attitude is above some (possibly contextually variable) threshold. I switch from talk of what attitude one "ought all-things-considered" to hold to talk of what attitude there is "decisive (or sufficient) all-things-considered reason" to hold for the purpose of achieving greater precision with respect to certain of the issues under discussion. Thanks to Selim Berker for helpful conversation.

First, a simple weighing procedure doesn't respect the different weighing behaviors exhibited by epistemic and practical reasons. If one has strong, equally-balanced epistemic reasons to believe p and to disbelieve p, and there are no other reasons (either epistemic or practical) in play, then, plausibly, one has decisive all-things-considered reason to suspend judgment on p. For example, suppose that a (trusty) weatherperson says that it'll be sunny, another says that it won't be, you have no other relevant epistemic reasons, and there are no practical reasons in play. Then, plausibly, you lack sufficient all-things-considered reason to believe that it'll be sunny, lack sufficient all-things-considered reason to disbelieve this proposition, and have decisive all-things-considered reason to suspend judgment. To generalize: if one has strong, equally-balanced epistemic reasons to believe p and to disbelieve p, and there are no other reasons on the scene, then one lacks sufficient all-things-considered reason to hold either attitude, and has decisive all-things-considered reason to suspend judgment.[10]

But if one has strong, equally-balanced practical reasons to believe p and to disbelieve p, and there are no other reasons in play, it's not at all plausible that one has decisive all-things-considered reason to suspend judgment on p. For example, suppose that I offer you $1 million to believe that the number of dust specks on my desk is even and, also, $1 million to disbelieve this. And suppose also that, aside from my pecuniary offers, you have no other reasons (either epistemic or practical) that bear on the matter. In this case, it would seem implausible to say that you have decisive all-things-considered reason to suspend judgment about whether the number of dust specks on my desk is even. Instead, it seems far more plausible to hold that you have sufficient all-things-considered

[10] Note that it's also very plausible that we get the same result even in circumstances in which the epistemic reasons for believing p and for disbelieving p are only *approximately* equal in strength.

reason to believe this proposition, have sufficient all-things-considered reason to disbelieve it, and lack sufficient all-things-considered reason to suspend judgment. To generalize: if one has strong, equally-balanced practical reasons to believe p and to disbelieve p, and there are no other reasons on the scene, then one has sufficient all-things-considered reason to hold either of these attitudes, and lacks sufficient all-things-considered reason to suspend judgment.[11]

But a simple weighing procedure can't respect this asymmetry in the weighing behavior of epistemic and practical reasons. This is because once the total reason *simpliciter* in favor of each doxastic alternative has been calculated, a simple weighing procedure will view as identical situations in which the only reasons in play are equally balanced in favor of believing p and disbelieving p, irrespective of whether the relevant reasons are epistemic or practical. In both kinds of situation, such a procedure will yield the result that one has equally strong reason *simpliciter* to believe p and to disbelieve p. So a simple weighing procedure is committed to delivering the same overall all-things-considered verdicts in both kinds of situation. But as we've just seen, we should get *different* verdicts, depending on whether the reasons in play are practical or epistemic. And so a simple weighing procedure doesn't respect the differing weighing behavior of epistemic and practical reasons for belief.[12]

---

11 This difference in the weighing behavior of epistemic and practical reasons for belief mirrors the well-known asymmetry in the weighing behavior of epistemic reasons for belief and practical reasons for action (see, e.g., Feldman [2000], Harman [2004], and Cohen [2016]). If one has equally strong practical reasons for two or more alternative actions, and there are no other reasons for action in play, then one has sufficient reason to perform any of the those actions, and lacks sufficient reason to perform any other. Likewise for practical reasons for belief. But not so with epistemic reasons for belief. Berker (MS) is the first and only author to date to point this out.

12 Berker (MS) raises a similar objection to the idea that what I'm calling a "simple weighing procedure" could be used to make comparisons among epistemic and practical reasons.

Second, as a mechanism for comparing epistemic and practical reasons, simple weighing is *too sensitive*. Suppose Sam will survive her illness only if she believes she will but that Sam's epistemic reasons for believing she won't survive are far weightier than those in favor of believing she will, such that Sam has decisive epistemic reason to believe she won't survive. And for the sake of argument, suppose Sam's practical reasons for believing she'll survive just barely outweigh her epistemic reasons for believing she won't survive, such that, by a hair's breadth, Sam has decisive all-things-considered reason to believe she'll survive. Now suppose there's a change in Sam's situation that generates a new (relatively weak) epistemic reason to believe she won't survive. Then the balance of reasons could easily tip toward its being the case that Sam has decisive all-things-considered reason to have this belief. But this is implausible. Given what's at stake, the addition of an epistemic reason for believing she won't survive can't plausibly make the difference in determining what Sam has decisive all-things-considered reason to believe. In fact, something stronger seems true: that in this case *no amount* of epistemic reason could oust the practical winner. The moral is that, in high-stakes cases, simple weighing is too sensitive to the total weight of epistemic reason.[13]

Each of these difficulties is on its own sufficient to show that a simple weighing procedure is inadequate for the purpose of comparing epistemic and practical reasons for belief. So each of these difficulties is instructive, in that it reveals a key desideratum that any satisfactory model for comparing epistemic and practical reasons needs to satisfy: (1) such a model must respect the differing weighing behaviors of the two types of reasons;

[13] On this, see Reisner (2008).

and (2) such a model should not be overly sensitive, in high-stakes cases, to the total weight of epistemic reason.[14]

3. Reisner's Proposal

In an article that offers the only existing contemporary account of how the two types of reasons might be compared, Andrew Reisner (2008) suggests a model for comparing epistemic and practical reasons that meets both these desiderata.[15] According to Reisner's proposal, if the weight of the practical reasons for or against some doxastic alternative is above some high threshold, then the epistemic reasons are irrelevant to determining what you have decisive or sufficient all-things-considered reason to believe,

[14] It's worth noting a third potential difficulty for the idea that epistemic and practical reasons for belief could be compared by a simple weighing procedure. A simple weighing procedure requires the assumption that the weights of the reasons being compared are commensurable, i.e., that there is some common unit in terms of which the weights of the relevant reasons can be measured (Chang 1997). This assumption is needed for the first step of the procedure in which the weights of the reasons for and against a given alternative are combined in order to determine the total amount of reason (*simpliciter*) in favor of that alternative. But the assumption that the weights of epistemic and practical reasons are commensurable is on its face implausible. For example, suppose we think the weight of an epistemic reason for believing p is determined by the degree to which the relevant reason indicates that believing p would promote (or somehow show respect for) the Jamesian goals of getting at the truth and avoiding error. And suppose we think the weight of a practical reason for believing p is determined by the degree of benefit that the relevant reason indicates that believing p would bring about. Then, plausibly, there will be no common unit in terms of which the weights of the two types of reasons can be measured. And if this is so, then simple weighing is a non-starter. I omit discussion of this issue in the main text since there are several ways one might skirt it. For example, one might try to give a unified account of the weights of epistemic and practical reasons (though this seems a tall order), or take facts about the weights of reasons to be absolutely basic (and simply assume that the weights of epistemic and practical reasons are commensurable in the relevant sense). No matter: either sort of view would face the difficulties raised in the main text.

[15] In the historical literature, a model for comparing epistemic and practical reasons is suggested by some remarks made by James (1897). James is often interpreted as having held that practical reasons can be relevant to determining what one ought-all-things-considered to believe only when the epistemic reasons are equally (or roughly equally) balanced in favor of belief and disbelief. In such cases, one's practical reasons are supposed to "break the tie." But while James's model might satisfy the first of our desiderata, it certainly doesn't satisfy the second. For example, in the case where Sam will survive her illness only if she believes she will, but Sam's epistemic reasons decisively favor believing she won't survive, James's model yields the (implausible) result that Sam ought all-things-considered to believe she won't survive. So James's model is too sensitive in high-stakes cases to the total weight of epistemic reason. So the model doesn't satisfy our second desideratum.

and the practical reasons alone determine the all-things-considered verdicts. Otherwise, the practical reasons are irrelevant, and the epistemic reasons alone determine the all-things-considered verdicts. In Reisner's lexicon: "when pragmatic reasons for belief are strong enough, [epistemic] reasons for belief are silent … otherwise, pragmatic reasons for belief are silent in determining what one ought to believe, all-things-considered" (2008, 24).

Two points of clarification. First, Reisner is explicitly neutral about whether the threshold that figures in his model is to be understood as having a fixed or variable value, and so he's open to a possible version of the proposal on which the relevant threshold is contextually determined (2008, 25-6). On the (orthogonal) matter of whether the threshold is vague or sharp, Reisner says nothing. But whichever way(s) we might want to go with regard to these cross-cutting distinctions, Reisner does think that on the most plausible version of his proposal, the threshold value will (tend to) be quite high, such that the epistemic reasons won't be easily overridden in all cases by the practical ones. More on why below.

Second, Reisner assumes that the weight of a practical reason for belief is a monotonically increasing function of the benefit that having the relevant belief would bring about (perhaps along with other variables).[16] So, roughly, a practical reason in favor of a belief is stronger, the better (i.e., more beneficial) having the belief would be (other things being equal). This is contentious, but I'll share with Reisner this assumption going forward.[17]

[16] Reisner's suggested account of the weights of practical reasons for *belief* thus resembles a value-based account of reasons for *action*. See Maguire (2016) for discussion and defense of the latter.

[17] I say that this is contentious only because one might think that there are some practical reasons for belief (e.g., certain moral ones) that aren't benefit related. Were this so, the above assumption concerning the

Reisner's model meets our desiderata. It's straightforward to see that, unlike a simple weighing procedure, Reisner's model isn't committed to yielding identical all-things-considered verdicts in situations in which the only reasons in play are equally balanced in favor of believing p and disbelieving p, irrespective of whether the relevant reasons are epistemic or practical. So Reisner's model can respect the differing weighing behavior of epistemic and practical reasons. So the model easily satisfies our first desideratum.

Moreover, Reisner's model avoids the "oversensitivity" worry that we had with a simple weighing procedure. On Reisner's model, once the strength of the practical reasons is sufficiently high, the epistemic reasons are irrelevant, and the practical reasons alone determine what you have decisive (or sufficient) all-things-considered reason to believe. So, once the practical reasons start to matter, no amount of epistemic reason can oust a practical winner. So Reisner's model readily satisfies our second desideratum, too.[18]

So Reisner's model succeeds where a simple weighing procedure fails. And besides this, the model has a further virtue. Suppose you're waiting in line to see a newly-released film and you overhear someone who's just seen the film reveal a major plot twist.[19] You've got excellent reason to believe the testifier is oblivious and inconsiderate, but none to think he's a liar. So your epistemic reasons strongly favor believing that the plot twist occurs. But clearly it would be better (i.e., more beneficial) to disbelieve this

---

weights of practical reasons would need to be rejected. However for the purpose of this paper, we can assume that all such apparently non-benefit-related practical reasons can be "consequentialized." This is only for simplicity's sake and won't affect the arguments to follow.

[18] Note also that Reisner's proposal doesn't require the assumption that the weights of epistemic and practical reasons are commensurable, and so it avoids the difficulty mentioned in note 14.

[19] I borrow this case from Kelly (2002).

proposition—believing it will ruin the film. So in this case, your practical reasons support believing that the plot twist doesn't occur. But intuitively, it's not the case that you have decisive (or sufficient) all-things-considered reason to have this belief. Instead, it seems that you have decisive all-things-considered reason to believe what your evidence suggests—that the plot twist occurs.

You might disagree. But I suspect this is just a matter of detail. The important point is that there are possible low-stakes cases in which it would be more beneficial overall to have an epistemically unsupported belief, but we think that, nonetheless, you ought to have the attitude that's best supported by the balance of your epistemic reasons. The further virtue of Reisner's model is that it has the resources to accommodate and explain this kind of judgment: assuming that the threshold value is set suitably high, the weight of the practical reasons in the relevant sorts of cases won't be strong enough to silence the epistemic ones. So, in such cases, the practical reasons will be irrelevant, and your epistemic reasons will alone determine what you have decisive or sufficient all-things-considered reason to believe.

So Reisner's model can accommodate the attractive idea that when the practical stakes are *not* particularly high, the practical reasons seem irrelevant to determining what you have decisive or sufficient all-things-considered reason to believe. Or to put the same point slightly differently: Reisner's model can make it the typical case that only one's epistemic reasons are relevant to determining all-things-considered doxastic verdicts. This is why Reisner thinks that whatever we might want to say about the nature of the threshold that his model posits (i.e., that it's fixed or contextually variable, vague or sharp), substantively speaking, it seems very plausible to think that the threshold should be set quite high.[20]

So Reisner's model has a lot going for it. But the proposal is ultimately inadequate. This is because Reisner's model faces a structural problem due to Selim Berker (MS). The problem is that, in certain circumstances in which the weight of your practical reasons is above the threshold but your practical reasons are equally balanced in favor of multiple doxastic alternatives, Reisner's model yields unacceptable results. For example, suppose that you have strong practical reasons against believing there's a terrible manticore in your midst at this very moment; call this proposition "q." You have excellent epistemic reasons to disbelieve q (I hope), but let's stipulate that the weight of your practical reasons against believing q is above the threshold, such that (according to Reisner's model) your epistemic reasons are irrelevant to determining the all-things-considered verdicts. Now, since all that your practical reasons tell you is "Don't believe q!," they leave it open whether you should disbelieve this proposition or suspend judgment on it. So Reisner's model yields the following all-things-considered verdict: you have sufficient all-things-considered reason to disbelieve q, have sufficient all-things-considered reason to suspend judgment on q, and lack sufficient all-things-considered reason to believe q. But this verdict is implausible. It's not at all plausible that

[20] This discussion reveals an important advantage of Inclusivism over Robust Pragmatism. On the most natural way of developing the latter view, it will deliver the result that, in every case, you ought to hold the attitude the having of which would bring about the most benefit. But, as we've just seen, this is intuitively implausible. And even if we supplement Robust Pragmatism with the further (admittedly plausible) thesis that it's in general beneficial to hold epistemically supported beliefs, there will still be *possible* low-stakes cases in which it would be slightly more beneficial overall to hold an epistemically unsupported belief, but we think that, despite this, you ought to have the attitude that's best supported by the balance of your epistemic reasons. And at least without some fancy footwork, it's hard to see how Robust Pragmatism could accommodate judgments of this kind. But, when paired with the right kind of model for comparing epistemic and practical reasons, Inclusivism clearly can accommodate such judgments, as the above discussion demonstrates.

in this situation you have sufficient all-things-considered reason to suspend judgment about whether there's a terrible manticore in your midst. Instead it's far more plausible to hold that you have decisive all-things-considered reason to *disbelieve* this proposition. But Reisner's proposal can't accommodate this judgment. And so the proposal is inadequate.

And the problem—Berker's—is perfectly general: In any situation in which (1) the weight of the practical reasons is above the threshold, (2) the practical reasons are equally balanced in favor of multiple doxastic alternatives, and (3) the balance of epistemic reasons favors one of those alternatives over the other(s), Reisner's model yields unacceptable results.

4. A New Model

What Berker's problem reveals is that in high-stakes cases in which the practical reasons are equally balanced in favor of multiple doxastic alternatives, the epistemic reasons can still seem somehow relevant to determining the all-things-considered verdicts. So, contrary to Reisner's proposal, while epistemic reasons can't plausibly unseat practical winners in high-stakes cases, these reasons do not, in such cases, become irrelevant. So what we need is a model that can accommodate this fact, but which also satisfies our initial two desiderata.

By introducing some additional machinery—in particular, the notion of lexical priority—I believe we can produce such a model.[21] According to the model that I suggest, if the weight of the practical reasons for or against some doxastic alternative is above

[21] The notion of "lexical priority" is sometimes intended to mean different things depending on the theorist deploying it, but I mean it in the way Rawls (its originator in moral philosophy) did.

some very high threshold, then the practical reasons become lexically prior to the epistemic ones, such that the attitude you have decisive all-things-considered reason to hold is the one that's best supported by the balance of epistemic reasons, among those that satisfy the prescription of the practical reasons. Otherwise, epistemic reasons are lexically prior to practical ones.[22]

This model yields all-things-considered verdicts identical to those yielded by Reisner's model in all circumstances except those in which the weight of the practical reasons is above the threshold, the practical reasons are equally balanced in favor of multiple doxastic alternatives, and the balance of epistemic reasons favors one of those alternatives over the other(s). To see this, suppose first that the threshold isn't met. Then my model predicts that the attitude that you have decisive all-things-considered reason to hold is the one that's best supported by the practical reasons among those that are compatible with the prescription of the epistemic reasons. But, since it's impossible for the balance of epistemic reasons to decisively (or sufficiently) favor two (or more) doxastic alternatives, there will be, in any given case, only *one* doxastic attitude that's compatible with the prescription of the epistemic reasons. So, in situations in which the threshold isn't met, this is the attitude that you have decisive all-things-considered reason to hold—and Reisner's model would agree.[23]

---

[22] Like Reisner, I'm agnostic about whether the threshold that figures in my model is best understood as being fixed or contextually variable, vague or sharp—but some discussion to follow (in section 5).

[23] The discussion in this paragraph assumes that epistemic permissivism is false. If epistemic permissivism is true, then there are possible cases in which the balance of epistemic reasons sufficiently favors multiple doxastic alternatives. And if such cases are possible, then the verdicts yielded by Reisner's model and mine will diverge in one more kind of case (other than that which I've already indicated), viz., cases in which the weight of the practical reasons for (or against) any of the relevant doxastic alternatives isn't above the threshold, the balance of epistemic reasons sufficiently favors multiple alternatives, and the practical reasons favor one of those alternatives over the other(s). In such a case, my model predicts that you have decisive all-things-considered reason to hold the attitude that's best supported by the practical reasons among those that are compatible with the prescription of the

Next, suppose that the weight of the practical reasons *is* above the threshold, but that the practical reasons aren't equally balanced in favor of two or more doxastic alternatives—that the preponderance of practical reason favors one doxastic option, in particular. Then my model predicts that the attitude that you have decisive all-things-considered reason to hold is the one that's best supported by the balance of epistemic reasons among those that are compatible with the prescription of the practical reasons. But since, by hypothesis, the balance of practical reasons favors a *single* doxastic alternative, there's only *one* attitude that's compatible with the prescription of those reasons. So that's the attitude that you have decisive all-things-considered reason to hold. And, again, Reisner's model would agree.[24]

Finally, suppose that the weight of the practical reasons is above the threshold, the practical reasons are equally balanced in favor of multiple doxastic alternatives, but the balance of epistemic reasons doesn't favor one of those alternatives over any other. (Incidentally, this is possible only when the weight of the practical reasons is above the

---

epistemic reasons (i.e., the most beneficial attitude among those that are sufficiently favored by the balance of epistemic reasons). Reisner's model, on the other hand, predicts that you have sufficient all-things-considered reason to hold either of the attitudes for which there is sufficient epistemic reason. Unsurprisingly I'm inclined to think that the verdict that my model delivers is preferable. Although I wish to remain officially neutral about whether epistemic permissivism is true, the remainder of the paper is written as if the view is false. Beyond what I've said in this footnote, however, nothing of importance turns on this.

[24] Suppose that the weight of the practical reasons is above the threshold, the preponderance of practical reason favors believing p, but suspending judgment on p is a close practical second. And suppose also that I have decisive epistemic reason to suspend judgment on p. In this situation, Reisner's model and mine (as I've stated it above) yield the verdict that I have decisive all-things-considered reason to believe p. I find this verdict acceptable, but some people report finding it intuitive that in this case I have at least sufficient all-things-considered reason to suspend judgment on p. To accommodate this judgment, one could simply specify that practical reasons for belief are to be compared in a satisficing way, such that one has sufficient practical reason for a given doxastic attitude just in case the total weight of practical reason in favor of the attitude is above a certain threshold. If suspending judgment on p is a close practical second to believing p, then if the total weight of practical reason in favor of believing p meets the relevant threshold, then, presumably, so too will the total weight of practical reason in favor of suspending judgment on p. And if so, then my model will yield the verdict that I have sufficient all-things-considered reason to hold either of these attitudes, and lack sufficient all-things-considered reason to disbelieve p. Again, I don't feel that this gambit is necessary. But in case one does, it's worth noting that the option is available.

threshold, there's exactly one attitude that's incompatible with the prescription of the practical reasons, and that same attitude is the one that's decisively or sufficiently favored by the balance of epistemic reasons.) Here, my model predicts that the attitude that you have decisive all-things-considered reason to hold is the one that's best supported by the balance of epistemic reasons among those that are compatible with the prescription of the practical reasons. But since, by hypothesis, the balance of epistemic reasons doesn't favor one of the alternatives compatible with the prescription of the practical reasons over any other, the doxastic attitude that you have decisive all-thing-considered reason to hold is one of *any* of those that are compatible with the prescription of the practical reasons. And Reisner's model agrees.

Now consider the situations that render Reisner's model inadequate—i.e., those in which the weight of the practical reasons is above the threshold, the practical reasons are equally balanced in favor of multiple doxastic alternatives, and the balance of epistemic reasons favors one of those alternatives over the other(s). So far, we've seen that in all other situations, Reisner's model and mine deliver identical verdicts. But in situations of this last sort, the verdicts yielded by Reisner's model and mine diverge. Consider again the case of the terrible manticore. Recall that in this case Reisner's model yields the verdict that you have sufficient all-things-considered reason to suspend judgment about whether there's a manticore in your midst. This seems wrong, given the details of the case. But my model does better. In the case of the manticore, the attitude that's best supported by the balance of your epistemic reasons, among those that satisfy the prescription of your practical reasons, is *disbelief* that there's a manticore in your midst. So on the model that I suggest, this is the attitude that you have decisive all-things-considered reason to hold. And my model yields similarly plausible results in all possible versions of situations of this sort.

So, unlike Reisner's model, my model avoids Berker's problem. Moreover, in much the same ways as Reisner's model, my model satisfies our initial two desiderata: it can respect the differing weighing behaviors of epistemic and practical reasons, and it predicts that, in high-stakes cases, no amount of epistemic reason can unseat a practical winner, thus avoiding the "oversensitivity" worry that we had with a simple weighing procedure. Finally, my model, too, can accommodate the highly attractive idea that when the practical stakes are *not* particularly high, the practical reasons seem entirely irrelevant to determining all-things-considered doxastic verdicts. Assuming that the threshold value is set suitably high, my model, like Reisner's, makes it the typical case that only your epistemic reasons will seem relevant to determining what you have decisive all-things-considered reason to believe.[25]

5. Arbitrariness

This section addresses a possible objection to my proposal. The potential worry begins with the following, natural question: how *much* must be at stake before the practical reasons become lexically prior to the epistemic ones? That is: how strong must the practical reasons for (or against) some doxastic alternative be, in order to meet the threshold my model posits?

As we've already seen, it seems very plausible that the threshold should be set quite high. And in order to avoid extensional equivalence with Alethism, the threshold

[25] My model achieves this result in a slightly different way than Reisner's. The "seems" qualifier is important: the practical reasons aren't *irrelevant*, but only lexically posterior in the relevant cases.

value of course needs to be finite. But this leaves an infinite range of possibilities. And wherever we might think the line should be drawn, one might ask: why *there*? Since the threshold could always be higher or lower, any particular location is likely to seem arbitrary. And this seems like a cost not only for my model, but also for the normative view for which it's designed—Inclusivism. After all, whatever else we might say about Alethism and Robust Pragmatism, these two views at least seem less arbitrary in their determinations of all-things-considered verdicts.[26]

I should note straight away that this worry can't easily be avoided by adopting a specific view about whether the threshold that my model posits is fixed or variable, vague or sharp. Perhaps the problem is at its most acute if we adopt a version of the proposal on which the threshold is fixed and sharp. But even on a version that says that the threshold is variable and vague, similar problems surface. Specifying an imprecise *range*, as opposed to a sharp cut-off, doesn't make the relevant range any less arbitrary. And holding that the threshold value is contextually determined, as opposed to fixed, doesn't help much either. For even if there were no fixed value for the threshold, its location in any given case instead being a function of the nature of the situation at hand, it seems far from obvious what the relevant "threshold function" should look like. In fact, it may be that no such function could be precisely specified. And if this is so, then determining the

[26] This point shouldn't be overstated. Both Alethists and Robust Pragmatists need to adequately specify plausible comparison functions for determining all-things-considered doxastic verdicts. And least for Alethists, worries about arbitrariness loom. This is because many Alethists favor the view that whether you ought to believe something isn't a matter of whether there is more or most epistemic reason for you to believe it, but rather of whether there is *sufficient* epistemic reason for you to believe it (in other words, these theorists opt for a satisficing way of comparing of the total amount of epistemic reason in favor of each of the relevant doxastic alternatives in order to determine all-things-considered verdicts—see note 9). But precisely how much epistemic reason must there be in favor of a given doxastic option in order for the epistemic reason in its favor to be sufficient? Worries about arbitrariness seem as potent here as they do for my proposal. (Of course, Alethists who aren't also satisficers won't face this worry; see Schroeder [2015] for discussion.)

correct cut-off point (or range) for the threshold would seem an unprincipled matter, requiring case-by-case judgment.[27]

So the charge of arbitrariness can't easily be avoided by adopting a particular view about the nature of the threshold that my model posits. What to say, then? I think the best thing to do is come clean—to admit that any location for the threshold will seem arbitrary, given that it could have always been higher or lower. The important question is whether this poses a serious difficulty, and I don't think that it does. Instead, I view the apparent element of arbitrariness in my model for comparing epistemic and practical reasons as a necessary feature not just of the particular normative view for which it's designed, but of *any* normative theory that posits two or more intrinsically significant normative factors, which ultimately need to be balanced against one another in order to determine all-things-considered verdicts.[28] If a normative theory posits multiple, intrinsically significant normative factors, and says that at least in certain circumstances

[27] One suggestion that doesn't particularly help us here, but which is nonetheless worthy of note, is that the location of the threshold might vary according to the "epistemic importance" of the proposition under consideration. Roughly: the more epistemically important the proposition, the higher the threshold. This is analogous to the common view among moderate (threshold) deontologists that the location of the threshold that their theory posits is in part a function of the nature of the particular deontological constraint to be infringed upon. In principle, I'm sympathetic to this suggestion. My only apprehension is that I find it difficult to grasp a notion of "epistemic importance" that doesn't collapse into *practica*l importance—in particular, the practical importance of having an epistemically supported attitude with respect to the relevant proposition. For a helpful discussion of this puzzle, and of what's needed to answer it, see Treanor (2014).

[28] What is it for a normative factor to be intrinsically normatively significant? I suggest the following (first pass) account: For a normative factor, F, to be intrinsically normatively significant is for F to be normatively significant just in virtue being the kind of factor that it is, and not in virtue of any relation it might bear to a normative factor of any other kind. Epistemic considerations are thus intrinsically normatively significant for belief, inasmuch as they're normatively significant just in virtue of being considerations of epistemic relevance, and not in virtue of, e.g., their being of practical import (if they are). This notion—that of intrinsic normative significance—provides a useful way of distinguishing Inclusivism from its competitors. For example, Robust Pragmatists (e.g., Rinard [2015]) claim that practical considerations are intrinsically normatively significant for belief, but deny that epistemic considerations are. Instead, these theorists claim that epistemic considerations are normatively significant for belief only when (and because) believing in accordance with them would provide some benefit. The Inclusivist denies this. Her core commitment is that both epistemic *and* practical considerations are of intrinsic normative significance for belief.

those factors need to be compared against one another, then that theory will require some sort of trade-off schedule—this is an inevitable consequence of normative pluralism of any kind.[29] And for any trade-off schedule that a pluralistic normative theory might posit, one might always ask: why *that* schedule, as opposed to some other? But this is to raise an epistemological question—viz., what evidence is there for the selected schedule rather than some alternative? And the answer to this question is likely to turn on a variety of factors—e.g., the degree to which the selected schedule can accommodate our intuitions across cases. But I see no good reason to think that such evidence can't help us home in on the location of the threshold that my model, in particular, posits. Indeed, notice that we've been collecting intuitions that provide us with some indication of where the relevant threshold might lie over the course of this paper.

So it's hard to see how the kind of pluralism about doxastic normativity that my model is meant to supplement could be any worse off than any other form of normative pluralism, at least with respect to the charge of arbitrariness. So unless we think that we should, for reasons of arbitrariness, abandon the possibility of pluralism *in general*, we shouldn't think the objection that we're considering speaks decisively against the proposal of this paper.

Of course, that's not to say that their relative lack of apparent arbitrariness doesn't provide Alethism and Robust Pragmatism with *some* advantage over Inclusivism and the particular model for comparing epistemic and practical reasons I've proposed for it. The simplicity and elegance of the former normative theories, particularly with regard to their determinations of all-things-considered verdicts, are theoretical virtues—and certain of

[29] Kagan (1998) makes this point. Much of the above material is inspired by Kagan's commentary concerning thresholds in his discussion of the prospects for moderate deontology (pp. 78-84.)

their proponents rightly tout them as such (Rinard 2017). But on this matter, I should like to echo the paradigmatic pluralist, W.D. Ross:

> [I]t is more important that our theory fit the facts than that it be simple, and the account we have given above corresponds (it seems to me) better than either of the simpler theories with what we really think. (1930, 19)

CHAPTER 4: Fitting Love and Reasons for Loving[1]

1. Introduction

Suppose Sam happens by a stray boot in the park and spontaneously comes to love it deeply and dearly. He wraps the boot in a blanket before placing it in the front seat of his car and, prior to driving off, makes sure it's buckled in safely. Sam tends to his beloved boot for years thereafter, relishing his time with it, and dreading any time he must spend in its absence.

We feel Sam's attitude is perverse. We want to say that his love for the boot is unintelligible, irrational, or crazy—that it's unfitting to feel the way that he does about it. We would feel the same way about someone in love with sour milk, Trump's literary style, or the number 246. This suggests that love is the sort of attitude that can be fitting or unfitting to its object. And this, in turn, raises the question: what sorts of considerations *make* love fitting?

This paper defends the following, simple answer: What makes love fitting are the lovable qualities of the beloved. Just as a person's admirable qualities make that person fitting to admire, and an outcome's desirable qualities make that outcome fitting to

[1] This paper has been preliminarily accepted for publication in the *Oxford Studies in Normative Ethics*, vol. 8 (Oxford University Press, forthcoming).

desire, the lovable qualities of the beloved are what make the beloved fitting to love. This is the so-called "quality view."[2]

The quality view promises a highly natural account of what kinds of considerations can make love fitting. Moreover, it has a number of theoretical virtues. For example, it offers a simple and elegant explanation of our reaction to Sam's bizarre love for his boot: Sam's love is unfitting because it's based either on lovable qualities his (unlovable) boot lacks, or on qualities of the boot that don't make it lovable.

But despite its attractions, the quality view is often quickly dismissed in the recent literature on love. This is because the view seems to face a battery of embarrassing difficulties. If you have lovable qualities that make it fitting for me to love you, does that mean that it's fitting me to love anyone who has those qualities? (The problem of promiscuity.) That if someone is more lovable than you, it's fitting for me to love her more? (The problem of trading up.) That if you lose your lovable qualities, it would be unfitting to love you? (The problem of inconstancy.) That it's fitting for anyone aware of your qualities to love you, too? (The problem of universality).[3]

And that's just a sample. In addition to the "big four" listed above, the quality view seems to face several other difficulties. But in this paper I'll address each of the putative problems facing the quality view. The defense is intended to be comprehensive.

[2] Defenders of the quality view, as I've characterized it, include Badhwar (1987), Delaney (1996), Velleman (1999), Keller (2000), Abramson and Leite (2011), Jollimore (2011), and Clausen (2016). The positions defended by these authors are unified by their commitment to the claim that personal, lovable qualities of the beloved are what make love fitting. They are diversified, in general, by their commitments concerning what specific qualities comprise *lovability*. For example, according to Velleman (1999), what makes someone lovable, and so fit to love, is her bare Kantian personhood, whereas for Abramson and Leite (2011), what makes someone fitting to love are the morally laudable non-relational features of her character.

[3] I borrow the names for these problems from Jollimore (2011).

One notable virtue of the defense of the quality view I'll provide is that it doesn't rely on a substantive account of the particular properties that comprise lovability. In this respect my defense differs from extant defenses of the quality view, all of which attempt to answer the various problems the view appears to face by advancing and defending highly specific accounts of the qualities that make people lovable.[4] Each such defense of the quality view is correspondingly committal in this regard. But my own defense is not. I'll remain neutral with respect to the issue of exactly what qualities comprise lovability.

2. Preliminaries

Love is complicated. I won't try to give an exhaustive account of its nature. Occasionally I will make a claim about love's essence, but when I do the claim will be uncontroversial. Is love an emotion (Brown 1987)? Does it constitutively involve a desire for union with the beloved (Nozick 1989)? Does it lie in the essence of love that if you love someone, then you take his interests as reasons to serve those interests (Frankfurt 2004)? I don't know. I won't take a stand. For now I assume only that love is properly understood as a distinctive kind of attitude that we might take toward persons or things.

Most parties to the debate concerning what kinds of considerations can make love fitting focus only on personal love—i.e., love that takes a person as its object.[5] These authors offer accounts of what makes it fitting to love romantic partners, family members, friends, colleagues, etc. But I won't limit myself in this way here. Although

---

[4] One exception is Clausen's (2016) defense of the quality view, which has some structural similarities to my own. I'll be noting some important differences between my view and Clausen's, as they become relevant.

[5] Some limit their focus even further to "reactive" love between adults (Abramson and Leite 2011), or to *romantic* love (Protasi 2016), or to a highly specific species of romantic love between adults (Bagley 2015).

most of my examples will concern personal love, the version of the quality view that I'll defend in this paper can easily be generalized to cover love for favored animals, football teams, films, and fiction. I take this to be a considerable advantage of the account.

A few words about what the quality view is and is not. The quality view, as I'll understand it, claims that the kinds of considerations that can make love fitting are specifically *intrinsic* qualities of the beloved, such as his kindness, wit, empathy, and grace.[6] Some authors explicitly or implicitly regard themselves (or are regarded by others) as quality theorists, but allow that extrinsic qualities of the beloved can make love fitting (see, e.g., Delaney 1996, Keller 2000, and Protasi 2016). But I think this is a mistake for two reasons. First, quality theorists who hold that only intrinsic qualities of the beloved can make love fitting can easily capture the natural and attractive thought that what makes love fitting must be something about the beloved *in his own right*. But quality theorists who say that extrinsic qualities can make love fitting can't similarly claim this virtue. Second, authors who identify as quality theorists, but claim that extrinsic qualities can make love fitting, seem to do so only because (they believe) this allowance helps them to solve the various problems the quality view appears to face.[7]

---

[6] I'm understanding "intrinsic qualities" in roughly the way Lewis (1983) suggests (though I don't pretend that theorizing about intrinsicality hasn't advanced since Lewis's work on the topic). Intrinsic qualities of the beloved are qualities she has in virtue of the way that she herself, and nothing else, is. A nice (but imperfect) test for intrinsicality is duplication. If I have an intrinsic quality, so does my perfect duplicate. An underappreciated fact in the love literature is that the intrinsic/extrinsic property distinction isn't coextensive with the distinction between non-relational and relational properties. All extrinsic properties are relational, but so are some intrinsic properties, e.g., the property of having a proper part. So the kind of quality view I'll defend in this paper doesn't rule out that certain relational qualities could make love fitting.

[7] For some straightforward (and recent) evidence for this sociological hypothesis, see Naar (2017).

But, as I'll argue, quality theorists needn't allow that extrinsic qualities of the beloved can make love fitting, in order to address the problems that seem to plague them.[8]

Now for what the quality view is not. The kind of quality view I'm concerned to defend does not imply that lovable qualities are the *objects* of love. When you tell your beloved that you love her on account of her remarkable poise, you aren't thereby claiming that the object of your love is your beloved's *poise*. Instead, what you're claiming is that you love *your beloved*, and that her remarkable poise partly grounds (or makes *pro tanto* fitting) that love. The version of the quality view I'll defend in this paper doesn't contradict this plausible diagnosis. The kind of quality view I'm interested in says that the lovable qualities of your beloved are what make your love for her *fitting*. It doesn't claim, implausibly, that these lovable qualities are the objects of your love.[9]

Next, some remarks about *fittingness*. The fittingness relation can be glossed as the relation in which a response stands to an object when the object merits—or is worthy of—that response. The quality view is thus the view that what makes somewhat merit love—what makes him worthy of it—are the lovable (intrinsic) qualities that he possesses. To characterize fittingness a bit further, we can note some connections it bears to certain other normative properties and relations. There's wide consensus that something is *valuable* just in case it's fitting to value, and that parallel equivalences hold between more specific value properties and the fittingness of certain, correspondingly specific ways of valuing. For example, it's very plausible that someone is *admirable* just

[8] Despite these considerations (and my personal, contrary preference), no conclusion I'll argue for in what follows contradicts the claim that certain extrinsic qualities can make love fitting. At the end of the day, I take it to be an interesting and otherwise worthy philosophical exercise to try to the defend the most thoroughgoing, unalloyed sort of quality view—the kind of quality view that says that only intrinsic qualities of the beloved can make love fitting. So that is the kind of quality view I will defend in this paper.

[9] On the distinction between object and ground invoked here, see Delaney (1996) and Keller (2000).

in case she's fitting to admire, and that something is *desirable* just in case it's fitting to desire. There's also some consensus that fittingness is a permissive rather than a requiring notion. You're not required to value everything that merits being valued (let alone to love everyone who's worthy of love), though it is plausible that, at least other things being equal, you're permitted to do so.[10] And lastly, it's widely held that considerations that make attitudes fitting provide (normative) reasons for those attitudes. For example, if the fact that you're intelligent makes you a fitting object of admiration, then, plausibly, that fact provides a reason to admire you. Similarly, if the fact that you're empathetic makes you fitting to love, then that fact plausibly provides a reason to love you. So there seems to be necessary covariance, at least in one direction, between the fittingness relation and the reason relation. But beyond these reasonably clear connections, the question of how fittingness metaphysically relates to other normative properties and relations lies fairly open. In a way, this will be our main topic of discussion.[11]

Finally, let me be clear that, aside from the motivation for the idea sketched at the outset of this paper, I'll provide no positive argument that love can properly be assessed as fitting or unfitting to its object. I'll simply be assuming that this is so in what follows. This assumption is somewhat controversial, though increasingly less so. Most authors who argue that love can't be fitting (or unfitting) appeal to either (or both) of two considerations: (1) that no existing account of what makes love fitting is satisfactory; and (2) that love is nonvoluntary (or, as it's sometimes more poetically put, "blind"), and so

[10] For discussion, see especially McHugh and Way (2016).

[11] For historical discussions of the notion of fittingness, see, among others, Brentano (1889/2009), Broad (1930), and Ewing (1948). For more recent discussions of the notion, see in particular D'Arms and Jacobson (2000), Chappell (2012), McHugh and Way (2016), Svavarsdottir (2014), and Howard (forthcoming [chapter 1 of this dissertation]).

isn't subject to the relevant kind of normative assessment.[12] But responses to the latter consideration have been rehearsed (*ad nauseum*) in the literature on love,[13] and the defense of the quality view I'll provide in this paper yields a counterexample to the former.

3. The Big Four

I'll start by addressing the "big four" objections to the quality view: the problems of promiscuity, trading up, inconstancy, and universality. I'll suggest a unified solution to all four problems, as opposed to piecemeal solutions to each. Then, a bit later (in section 4), I'll draw on certain components of my proposed solution to the "big four," in order to answer the remaining difficulties purportedly faced by the quality view.

Let's begin by rehashing the "big four." The quality view claims that what makes love fitting are the lovable qualities of the beloved. This seems to imply (1) that if your beloved has lovable qualities that make her fitting to love, then it's fitting for you to love anyone who has those qualities (the problem of promiscuity); (2) that if someone is more lovable than your beloved, it's fitting for you to love him more (the problem of trading up); (3) that if your beloved loses her lovable qualities, it would be unfitting for you to love her (the problem of inconstancy); and (4) that it's fitting for anyone who's aware of your beloved's lovable qualities to love him, too (the problem of universality).

My response is straightforward: I accept all of these implications. I accept, for example, that if newcomer Max really is more lovable than your beloved partner George, then it *is* fitting for you love Max more than George; that if George loses all of his

[12] See, for example, Frankfurt (1999, 2006), Zangwill (2013), and Smuts (MS).
[13] One of the earliest responses is due to Kolodny (2003); the most recent is due to Naar (2017).

lovable qualities, it *would* be unfitting for you to love him; and that if George has lovable qualities that make him fitting to love, then it would be equally fitting for you to love absolutely anyone else who happens to have those very same qualities. But what I deny, crucially, is that from the fact that it would be *fitting* (unfitting) for you to love someone, it necessarily follows that you have sufficient reason (not) to do so. To claim otherwise would be to fall victim to what Justin D'Arms and Daniel Jacobson (2000) call the "conflation problem": the problem of conflating the fact that it's fitting to feel a certain way with the fact that there is sufficient (or decisive) reason for that feeling.[14] I submit that this move mitigates the force of the "big four" considerably if not entirely.[15] The quality theorist can happily accept, e.g., that if someone is more lovable than your beloved, then it would be fitting for you to love that person more, without being committed to the implausible claim that you thereby have sufficient all-things-considered reason to do so.

So quality theorists should deny that facts about which attitudes are fitting necessarily amount to facts about which attitudes there's sufficient reason to have. Is this plausible? Yes. Although rejecting the relevant equivalence between *fittingness* and *reasons* isn't uncontroversial, the position has many proponents (e.g., Rabinowicz and Rønnow-Rasmussen 2004, Danielsson and Olson 2007, Rosen 2015, Howard

[14] D'Arms and Jacobson introduce and discuss the conflation problem in a different context. They present the problem as a problem for certain forms of neo-sentimentalism. In particular they argue that existing neo-sentimentalist theories lack the resources to distinguish moral and prudential assessments of certain sentiments from assessments of the appropriateness, or fittingness, of those sentiments. According to D'Arms and Jacobson, this poses a threat to the sentimentalist idea that values are grounded in the sentiments: very roughly, if sentimentalists can't draw the relevant distinction, this undermines their account(s) of value. But as many have noted, the conflation problem isn't a problem only for sentimentalists—it's relevant to a variety of other debates. For discussion of the generality of the problem, see Schroeder (2010).

[15] Later (in sect. 6), I discuss the worry that the move in question doesn't deal entirely with the "big four."

forthcoming).[16] And the work the distinction does for quality theorists, in particular, suggests that they should follow suit. And notice that the move in question doesn't commit the quality theorist to denying the highly plausible claim that considerations that make love fitting provide genuine (normative) reasons for love. The quality theorist can maintain that the lovable qualities of people provide genuine "fit-related reasons" for loving them; she just has to hold that fit-related reasons for love aren't the *only* reasons for love that there are.

What other sorts of considerations might provide reasons for love? Notice that each of the purportedly implausible implications of the quality view we've considered so far seems implausible only if we assume the presence of a valuable, loving relationship. Take the problem of trading up. If you're not in a valuable, loving relationship with George, then why *shouldn't* you love Max more than George, if Max is more lovable? And this point generalizes to all the "big four" problems facing the quality view. Consider the problem of promiscuity. If Maxine has lovable qualities that make her fitting to love, and you're not already in a valuable relationship with someone else, then why shouldn't you love Maxine, or anyone else who shares her lovable features? Or take the problem of universality. If you're not in a valuable relationship with Max, then why shouldn't anyone who's similarly unattached and aware of his lovable qualities love him, too? And consider finally the problem of inconstancy. If Maxine is truly unlovable, and

[16] For two prominent, recent defenses of the equivalence in question, see Shah (2006) and Way (2012). For my response to the former piece, see Howard (2016a; chapter 2 of this dissertation); for my response to the latter, see Howard (2016b).

you're not already in a valuable, loving relationship with her, then, barring reasons of beneficence, why think that you have any reason to love Maxine at all?[17]

What all this reveals, I think, is that valuable relationships have normative significance for love.[18] In particular, since loving someone with whom you're in a valuable relationship is necessary for the continued existence of that relationship, the presence of the relationship provides derivative "value-related" reasons for love. For example, in cases where, intuitively, I shouldn't trade up, what explains this is the presence of value-related reasons, which are (ultimately) provided by the value of my extant relationship. Similarly, in a situation in which my beloved loses her lovable qualities, but it seems that I shouldn't cease loving her, this is due to my value-related reasons to continue loving her, which derive from our valuable relationship. And similar appeals to the presence of value-related reasons provided by valuable, loving relationships are, I think, sufficient to adequately explain the rest of the relevant data.

The picture that emerges is this. The considerations that make love fitting are the lovable qualities of the beloved. These qualities provide genuine (fit-related) reasons for love. But the fit-related reasons provided by a person's lovable qualities aren't necessarily the only reasons to love him. If you're in a valuable, loving relationship with the relevant person, then, in addition to the fit-related reasons to love him, you have value-related reasons to love him, which derive from the value of your valuable, loving

---

[17] On some versions of the quality view, it would impossible for someone to lack qualities that make her lovable. For example, Velleman's (1999) claim that any given person is lovable in virtue of her bare Kantian personhood would seem to imply this result. Since I want to remain agnostic about what qualities make people lovable, I also want to remain neutral about whether it's possible for someone to be unlovable.

[18] That valuable relationships have normative significance for love shouldn't be all that surprising. A number of writers have suggested as much (e.g., Hurka 1997; Kolodny 2003). But the specific account of the normative significance of valuable relationships for love I'll proceed to offer is, in fact, a new suggestion.

relationship. This is the picture of fitting love and reasons for loving I think the quality theorist should adopt.[19]

4. The Relationship View

On the picture I've suggested, the normative significance of a valuable relationship for love consists in its being a source of value-related reasons for love. Valuable relationships are normatively significant for love, on my account, because they contribute to explanations of why your continued love for your beloved would be somehow valuable, or *good*. But Niko Kolodny (2003) offers a different account of the

---

[19] Clausen (2016) imparts a picture that's in some respects similar. In particular she suggests that lovable qualities of the beloved are what make love fitting, and that valuable relationships provide reasons (not) to act in ways that (threaten) facilitate the existence of those relationships. So on Clausen's view, if someone is more lovable than your beloved, it's fitting for you to love that person more, but it doesn't necessarily follow that you have decisive (or sufficient) reason to behave toward the relevant person in any particular way. Thus Clausen writes: "[O]ne might have strong fittingness reasons to love someone, but reason not to act in the ways that love would incline one to act in other circumstances, because a prior relationship prohibits such behavior" (p. 18). Further, Clausen suggests that "in cases where pursuing a certain kind of relationship with a fitting object of love is prohibited by the norms of an extant relationship, what would otherwise count as a reason for pursuing such a relationship would in this case be silenced" (p. 18). Now perhaps the most important difference between the picture that I've suggested and Clausen's is that on my view, value-related reasons provided by valuable loving relationships favor continued *love*, as opposed to merely favoring the continued performance of the various types of *action* that characterize one's participation in the relevant type of relationship. I take this to be an advantage of my account over Clausen's. It wouldn't be much consolation to hear from your beloved that although she (fittingly) loves some newcomer more than she loves you, she'll refrain from *acting on* (but not feeling) intense love for her new beloved. Similarly, if your beloved lost his lovable qualities, it would be of little comfort for him to hear that, while you no longer love him (you have no reason to, after all), you will continue to act in ways that suggest that you do, for the sake of your shared history. In short: our beloveds want us to *love* them; not merely to act in ways that people tend to act, when they're trying to preserve or promote their relationships. My picture accommodates this judgment; Clausen's does not. Further, Clausen's suggestion that reasons to pursue a new relationship are *silenced* in cases where your current relationship "prohibits" such behavior seems too strong to be plausible. The prospective value of a loving relationship other than the one that you're currently in plausibly provides you with *some* reason to pursue it—it's just that, in a (large) class of cases, these reasons are outweighed; I develop this idea further in section 6. And finally, Clausen's account doesn't obviously provide the resources needed to solve the remaining problems for the quality view. While Clausen does suggest some piecemeal solutions to the problems that (purportedly) face the quality view beyond the "big four", she makes no attempt to draw on her picture of fitting love and reasons for loving, in particular, in order to address these remaining difficulties. But in section 5, I'll show that, in contrast, the picture I've proposed does in fact provide quality theorists with the relevant resources.

normative significance of valuable relationships for love. He claims that the fact that you have a valuable relationship with your beloved is what makes your love for her *fitting*. This is the "relationship view" of what makes love fitting.[20] It is the quality view's main competitor. But in this section I'll argue that the relationship view is false: the fact that your beloved is in a valuable relationship with you does not (and indeed, cannot) make your love for her fitting.

My argument against the relationship view is simple. Part of what it is to love someone is to value him noninstrumentally, for his own sake. If you don't value someone noninstrumentally, then you don't love him at all. So whatever makes love fitting must also make fitting noninstrumental valuation of the beloved. But a consideration can make it fitting to value someone (or something) noninstrumentally only if that consideration makes that person (or thing) noninstrumentally valuable. And while there are likely many facts about your beloved in virtue of which she's valuable for her own sake, the fact that you have a valuable relationship with her is not plausibly among them. So the fact that your beloved is in a valuable relationship with you doesn't make it fitting to value her for her own sake. It follows that this fact can't make your love for your beloved fitting.[21]

---

[20] As noted above (see note 17), Kolodny isn't the only defender of the relationship view. I focus mostly on Kolodny, and his version of the view, only because he's offered the most sustained defense of it. That being said, the argument that I provide in this section generalizes to apply to any version of the relationship view.

[21] A number of authors have argued that a valuable loving relationship isn't necessary for fitting love, focusing on cases of fitting love at first sight, and fitting unrequited love (see, e.g., Setiya 2014, Protasi 2016, and Clausen 2016). The arguments offered by these authors seem to me to be successful, but the argument I've just offered above is in an important way stronger, since it purports to show that a valuable loving relationship isn't *sufficient* for fitting love. This is a strictly stronger conclusion, because the claim that a valuable loving relationship isn't necessary for fitting love is compatible with the possibility that fitting love is "multiply realizable"—that a person's qualities *or* a valuable relationship with that person, could make love for that person fitting. The claim that a valuable loving relationship isn't sufficient for fitting love, however, entirely rules out the possibility that such a relationship could (ever) make love fitting.

I take it to be obvious that loving someone essentially involves valuing her for her own sake and that it follows from this that whatever makes love fitting must also make fitting noninstrumental valuation of the beloved. These claims seem uncontroversial. I also take it to be obvious that a given consideration can make it fitting to value someone (or something) noninstrumentally only if that consideration makes that person (or thing) noninstrumentally valuable. This much follows from the fact that *being (noninstrumentally) valuable* and *being fitting to value (noninstrumentally)* are equivalent. The above argument, then, turns entirely on the substantive, axiological premise that your beloved isn't made noninstrumentally valuable by the fact that you have a valuable relationship with her. Two clarifying remarks regarding this claim are worth making. First, in advancing the claim in question, I'm not suggesting your beloved *lacks* noninstrumental value, and thus that it's unfitting to value him for his own sake. Instead, I'm claiming only that the particular fact that you have a valuable relationship with your beloved isn't among the facts that *make* him noninstrumentally valuable, and so fitting to value as such.

Second, in claiming that the relevant relational fact about your beloved isn't plausibly one in virtue of which he's of noninstrumental value, I'm not presupposing that it's impossible for noninstrumental value to be grounded in a relational fact. This would be (close to) question-begging against the relationship view, given the plausibility of the essentialist claim that loving someone constitutively involves valuing her for her own sake.[22] Moreover, while I'm not particularly persuaded by the proposal, a number of

[22] Another would-be question-begging assumption is that loving someone (or something) essentially involves valuing that person (or thing) not just noninstrumentally, but *intrinsically*, in her (its) own right. As it happens, I find this idea very plausible. But since it's clearly question-begging, I won't rely on it here.

authors at least tentatively endorse the idea that relational facts do sometimes ground noninstrumental value (Beardsley 1965, Korsgaard 1983, Kagan 1998, Rabinowicz and Rønnow-Rasmussen 1999).[23] For example, Shelly Kagan suggests that the pen that Abraham Lincoln used to sign the Emancipation Proclamation might be valuable for its own sake in virtue of the (relational) fact that it was used by Lincoln to sign the Emancipation Proclamation. And Kolodny (2003) himself suggests that human remains have noninstrumental value in virtue of the relation they bear to the human whose remains they are. These claims may seem plausible. But my suggestion is that it's not similarly plausible that your beloved has noninstrumental value in virtue of the particular relational fact that you have a valuable relationship with her. So I'm not necessarily opposed *in principle* to the idea that relational facts can ground noninstrumental value; instead, I'm (officially) opposed only to this specific instance of the idea's application.

My claim that your beloved's being in a valuable relationship with you isn't a consideration in virtue of which he's valuable for his own sake is, I think, intuitively compelling. However one might think that at least Kolodny's version of the relationship view has the resources to resist it. To see why, we need to get a little more of Kolodny's view on the table. So consider first that, according to Kolodny, loving relationships have noninstrumental value in addition to any purely instrumental value they might provide.[24]

---

[23] For a recent argument that relational facts can't ground noninstrumental value, see Tucker (2016). And for some further resistance to the idea, see Zimmerman (2011). Note that if either of these authors are correct—if relational facts can't ground noninstrumental value—then the relationship view is a non-starter.

[24] Specifically, Kolodny thinks that certain types of loving relationships have *final value*, where something is finally valuable just in case it's both noninstrumentally and intrinsically valuable. If certain types of relationships really do have this kind of value, then it would be fitting to value those types of relationships noninstrumentally. But note that it doesn't follow from the fact that it's fitting to value *your relationship* noninstrumentally, that it would be fitting to value *your beloved* noninstrumentally. At any rate, see Keller (2013) for some criticism of Kolodny's suggestion that certain types of relationships have final value.

Moreover, loving relationships are individuated, in part, by their participants: your loving relationship with your beloved wouldn't be the particular relationship that it is if either you or your beloved didn't figure in it. So on Kolodny's picture, the fact that your beloved is in a valuable loving relationship with you amounts to the fact that she's a constitutive part of a noninstrumentally valuable whole, viz., your loving relationship. And perhaps one might think that this relational fact about your beloved does indeed make her noninstrumentally valuable, since one might suspect that, as a fully general matter, the fact that something is a constitutive part of a noninstrumentally valuable whole suffices to make that thing itself noninstrumentally valuable.[25]

But this suspicion would be mistaken. Something's being a constitutive part of a noninstrumentally valuable whole does not, thereby, make that thing noninstrumentally valuable. To borrow a nice example from Nicholas Stang (2012), consider a Beethoven piano sonata with a particularly dramatic rest. The rest is a constitutive part of the sonata, and the sonata (let's assume) is a noninstrumentally valuable whole. But the rest itself has no noninstrumental value—it's not fitting to value for its own sake. Instead the rest is fitting to value only in virtue of its constitutive contribution to a larger whole—the sonata—which is itself fitting to value for its own sake. And as a fully general matter: if something is fitting to value in virtue of its constitutive contribution to a (noninstrumentally) valuable whole, that thing isn't fitting to value *for its own sake* in

[25] This is the best I can do to construct a possible rationale for thinking that being in a finally valuable relationship with someone makes it fitting to value that person for her own sake. Kolodny seems to suggest that it's fitting for you to love your beloved (and thus to value her for her own sake) in virtue of the relation that she (necessarily) bears to your finally valuable relationship. Unfortunately he never specifies what this relation is, and I can't think of another candidate other than that of *being a constitutive part*. Since Kolodny thinks relationships are individuated in part by their participants, I don't think he'd disagree that your beloved *does* necessarily bear this relation to your relationship. So I suspect Kolodny's thinking is that it's fitting to love your beloved (and so to value her for her own sake) in virtue of the fact that she's a constitutive part of your finally valuable relationship. But as I'll go on to argue, this is ultimately misguided.

virtue of its contribution. So the fact that something is a constitutive part of a noninstrumentally valuable whole doesn't plausibly suffice to make that thing, itself, noninstrumentally valuable. So contrary to the suggestion we're currently considering, the fact that your beloved is a constitutive part of your noninstrumentally valuable relationship is not plausibly a fact in virtue of which she's of noninstrumental value.

I can think of only one other way in which relationship theorists might try to defend the idea that your beloved's being in a valuable relationship with you is among the considerations that make him noninstrumentally valuable. They might admit that the relevant relational fact about your beloved is not one in virtue of which he's of *agent-neutral* noninstrumental value, but maintain that it is one in virtue of which he's of *agent-relative* noninstrumental value. In particular, relationship theorists might propose that the fact that you have a valuable relationship with your beloved makes him noninstrumentally valuable *relative-to-you*. They might then suggest that once the relevant relativization is made explicit, the axiological claim under consideration seems far more plausible.

As far as I can make sense of the idea that someone (or something) has noninstrumental value "relative-to-someone," I admit that the claim that your beloved is noninstrumentally valuable relative-to-you in virtue of the fact that she's in a valuable relationship with you is in one respect more plausible than the non-relativized version of the claim we've been considering so far—viz., it avoids the extremely odd implication that the relevant relational fact about your beloved is one in virtue of which it would be fitting for *anyone* to value her for her own sake.[26] But beyond this I see no further advantage of the suggested appeal to "agent-relative value." As many authors have noted, the notion of agent-relative value is of questionably intelligibility.[27] Indeed, as Mark Schroeder (2007) suggests, it's not at all clear what it means to say that something (or someone) is "valuable-relative-to-you," if this is meant to be distinct from saying, for example, that that thing is valuable *for* you, or something you believe to be (agent-neutrally) valuable.[28] So inasmuch as we have little (to no) idea what it means to say that something has agent-relative value, it would seem difficult (if not impossible) to evaluate the relative plausibility of the proposal that your beloved is "noninstrumentally valuable relative-to-you" in virtue of the fact that she has a valuable relationship with you. And at any rate, insofar as agent-relative value looks philosophically dubious, the relationship view's commitment to its existence would make the view correspondingly less attractive. And notice that neither the account of what makes love fitting that I've offered above nor the supplementary account of the normative significance of valuable relationships with which I've proposed to pair it comes with this commitment. Indeed, no component of the

---

[26] This further worry for the idea that your beloved's being in a valuable relationship with you is among the considerations that make him noninstrumentally valuable assumes, plausibly, that something's being agent neutrally-valuable is (extensionally) equivalent to its being (*pro tanto*) fitting for anyone to value.

[27] See, among others, Dancy (2000), Schroeder (2007), Zimmerman (2011), and Maguire (2016).

[28] One might think that my suggestion that something's being agent-neutrally valuable is equivalent to its being fitting for anyone to value hints at a simple way of understanding agent-relative value (see above and note 25). If agent-neutrally valuable things are things that are fitting for *anyone* to value, then maybe agent-relatively valuable things are things that are fitting for *someone*, but not just anyone, to value. However, as Mark Schroeder (2007) has forcefully argued, this way of drawing the distinction can't be right. Suppose that some outcome, O, is agent-neutrally very valuable, but that O is disvaluable "relative-to-you." If the notions of agent-neutral value and agent-relative value are distinct, then this situation should be possible. But on the way of drawing the distinction we're currently considering, it's not. For on the proposal under consideration, if an outcome is agent-neutrally valuable, then it's fitting for anyone, including you, to value. So in the situation described, it's fitting for you to value O. But by hypothesis, O is *disvaluable* relative-to-you, and so not an outcome that's fitting for you to value. So this way of drawing the distinction can't be right. A different way of drawing the distinction is proposed by Suikkanen (2009). Suikkanen's suggestion avoids Schroeder's problem, but faces other difficulties, due to Zimmerman (2011). It's fair to say that no proposal about how to understand the notion of agent-relative value has gained wide acceptance.

picture that I've suggested requires commitment to the existence of agent-relative value, or to any other, similarly suspect "philosopher's trinket." So, even if relationship theorists opt for the response under consideration, they do so only at the cost of accruing significant theoretical disadvantage. And given the current dialectical situation, significant theoretical disadvantage isn't something that relationship theorists can plausibly afford to accrue since the principal motivation for their view over the quality view lies in its purported ability to avoid the (putative) difficulties faced by the latter; difficulties that, as I'm currently in the process of arguing, the quality view does not, in fact, face.

This concludes my defense of the claim that your beloved's being in a valuable relationship with you isn't among the considerations that make her noninstrumentally valuable. I can think of no other way for relationship theorists to resist this claim beyond those discussed. Since whatever makes love fitting must also make fitting noninstrumental valuation of the beloved, and since a fact can't make it fitting to value someone noninstrumentally unless that fact makes her noninstrumentally valuable, it follows from the claim in question that the fact that you have a valuable relationship with your beloved does not (and indeed, cannot) make your love for her fitting. The relationship view is false.

5. Modes of Love and the Amenesiac Biographer

I've just argued that the relationship view offers an unsatisfactory account of the normative significance of valuable relationships for love, since it claims that the normative significance of a valuable relationship for love consists in its making love *fitting*, and valuable relationships do not (and cannot) make love fitting. And just prior to that, I argued that by accepting a certain picture of fitting love and reasons for loving, quality theorists can easily avoid the "big four" problems standardly raised against them. But two (putative) problems for the quality view remain: the problems of "modes" and "amnesia" (Kolodny 2003). So in this section, I'll draw on certain components of my solution to the "big four" in order to address these remaining difficulties.

Consider first the problem of modes. You probably do—and more importantly for our purposes, *ought to*—love your romantic partner in a way that's very different from the manner in which, e.g., his mother ought to love him. For the quality theorist, the problem of modes is the problem of explaining why this should be. If your partner's qualities are what make it fitting for both you and his mother to love him, then why should the particular ways in which you and his mother love your partner differ so considerably?

The picture of fitting love and reasons for loving I've proposed suggests a simple answer. What explains the fact that you ought to love your partner in a way that's very different from the manner in which his mother ought to love him has nothing to do with the fit-related reasons provided by your partner's lovable qualities, but rather everything to do with the value-related reasons provided by the value of your loving relationship. On the account that I've suggested, you have value-related reasons to love your partner because your continued love for him is necessary for the continued existence of your valuable, loving relationship. And the same goes for your partner's mother. But crucially, the value-related reasons provided by your respective relationships don't favor loving your partner in precisely the same *ways*. This is for the simple reason that the kind of love that would facilitate the continued existence of your relationship with your romantic partner is very different from the kind of love that would facilitate the continued existence of his relationship with his mother. And so the value-related reasons provided by your respective relationships favor loving your partner in rather *different* ways. And this is why, despite the fact that your partner's lovable qualities are what make it fitting for both you and his mother to love him, the particular kinds of love that you and your partner's mother ought to feel for your partner differ so dramatically. So quality theorists who accept the picture of fitting love and reasons for loving that I've proposed have a simple and intuitive answer to the problem of modes.

Consider next the problem of amnesia. The problem of amnesia, as Kolodny characterizes it, is the problem of explaining why the loss of certain memories extinguishes love. Assuming (as we have been) that love is an attitude we can hold for reasons, a natural explanation of this phenomenon is, as Kolodny suggests, that the amnesia victim loses cognitive access to his reasons for love. And if this is right, then, according to Kolodny, quality theorists would need to explain amnesia's effect on love by appeal to its effect on the lover's awareness of the beloved's qualities. But in some cases, such an explanation seems unavailable. For instance, suppose that Tom is somehow highly knowledgeable about someone's personal qualities despite the fact that he's never met her—perhaps, to borrow Kolodny's example, Tom is a biographer who's produced a "strikingly intimate portrait" of his subject's life and character via the detailed testimonies of her family and friends (2003, 141). And suppose also that, eventually, Tom meets the person whose qualities he knows so well, that the two of them fall in love, and marry. But years later, Tom suffers severe memory loss. He can recall everything that happened to him up until a year before his relationship with his wife began, but his memories of everything that happened after are lost—Tom's forgotten his relationship with his wife. In this case, Kolodny claims that "we would not expect [Tom] to love [his wife], and indeed it is hard to see how he could" and that "[t]his is so even though [Tom's] beliefs about [his wife's] personal qualities, and his confidence in them, have changed only slightly, if at all" (2003, 141). So, Kolodny concludes, in this case, amnesia's effect on love can't be explained by appeal to its effect on the lover's awareness of the beloved's qualities. So for quality theorists, the problem of amnesia is the problem of explaining why, in the kind of case just described—viz., a case in which the lover forgets his relationship with the beloved, but not the beloved's qualities—the lover's love disappears.

My response to the problem of amnesia proceeds in two steps. First, I think quality theorists should simply reject the demand to explain why forgetting a relationship with someone (whose qualities you continue to know) would extinguish your love for that person, since that seems like no fact at all. It is surely possible to love someone without believing you have (or had) a relationship with her. This happens all the time.[29] So I think quality theorists should reject Kolodny's claim that, despite remaining fully aware of your beloved's qualities, forgetting your relationship would extinguish your love.[30]

[29] As I mentioned earlier, it seems very plausible that a valuable loving relationship isn't necessary for fitting love (see note 20). And if it isn't, then, *a fortiori*, such a relationship isn't necessary for love, either.

[30] I'm not denying that it's *possible* that your love for your beloved would disappear if you forgot your relationship with her. But this (mere) possibility doesn't seem like a philosophically interesting explanandum.

Still, something close to Kolodny's claim does seem plausible. For while it's not certain nor even necessarily likely that your love for your beloved would *disappear* in the kind of case we're considering, it does seem reasonable to expect that the character or intensity of the love that you feel for your beloved would *change*. Indeed, we can easily imagine Tom, our amnesiac biographer, now loving his wife in the way that he did prior to their relationship, or in its early stages, rather than in the way that he came to love her after years of marriage. But if this is the datum that quality theorists need to be able to explain, then those who accept the picture of fitting love and reasons for loving that I've suggested can easily deliver—and this is the second step of my response. For as I just argued, value-related reasons provided by valuable relationships favor loving people in particular ways, viz., in ways that facilitate the continued existence of the relevant relationship. So, prior to Tom's memory loss, we can assume that he possessed value-related reasons provided by his valuable relationship to love his wife in a particular way, viz., in a way that would help to ensure that their relationship persists. But, as an effect of the memory loss, Sam *lost* those reasons, in the sense that the relevant reasons are no longer in his epistemic ken. And this, I think, easily suffices to explain why, as a result of his memory loss, the kind of love that Sam feels for his wife is different from the kind of love that he felt before: from Sam's epistemic perspective, there simply are no considerations that favor loving the woman before his eyes *as his wife*. And so he doesn't. So quality theorists who accept the picture of fitting love and reasons for loving that I've suggested have a natural and attractive answer to the problem of amnesia, too.

6. Still Weird and Trading Up Redux

I've now given my answers to each of the problems the quality view is often claimed to face. I know of no other challenges to the view in the literature.[31] But all of my answers to the problems putatively faced by the quality view turn on the specific picture of fitting love and reasons for loving I've proposed for it. So I'll conclude by considering some objections to that picture: what I call the "still weird" objection and the problem of trading up redux.

6.1. *Still Weird*

In response to the "big four," I argued that quality theorists should deny that facts about which attitudes are fitting necessarily amount to facts about which attitudes there's sufficient reason to have. Rejecting this equivalence allows quality theorists to deny that from the fact that it would be fitting for you to love someone, it necessarily follows that you have sufficient reason to love that person. And I claimed that this gambit mitigates the force of the "big four" considerably if not entirely. But isn't it *still weird* to think that it could be fitting for you to love someone else's partner (or child!) more than your own? And doesn't it sound *odd* to suggest that it would be unfitting for you to love someone whom you've loved for a lifetime, if she lost the qualities that make her lovable, perhaps through no fault of her own? Questions like these constitute the core of the "still weird"

[31] There do of course remain challenges to specific *versions* of the quality view, i.e., to quality views that specify exactly *what* personal qualities comprise lovability. But my aim was never to defend a specific proposal about what qualities make people lovable. Instead my aim was to address each of the extant challenges to the claim that unites *all* versions of the quality view, viz., that what makes love fitting are lovable, personal qualities. And as I say above, I know of no further challenges to *this* claim in the literature.

objection. They rhetorically suggest that denying an equivalence between *fittingness* and *reasons* does not, in fact, do much to mitigate the force of the "big four."

In response, I think that once we're equipped with a proper understanding of the notion of fittingness, none of the purportedly peculiar implications of the kind of quality view I've been defending really do seem all that "weird." 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. So to say that it would be fitting for you to love someone else more than your beloved is to say only that that person merits or is worthy of more love than your beloved—i.e., that the relevant person is *more lovable*. And it would seem naïve and romantic (in a bad way) to deny the possibility that someone *could* be more lovable than your beloved—to believe that your beloved is the most lovable person there is. And the beauty, of course, is that he doesn't have to be; your beloved has other things going for him that suffice to make it the case that you have sufficient reason to love him more than any (even more lovable) alternative(s)—viz., he figures essentially in your valuable, loving relationship. And for broadly similar reasons, I don't find it problematic to hold that if your beloved loses his lovable qualities, it would be unfitting to love him. Indeed, one of the most wonderful things about the sort of love that we feel for those we love most is that it can (rightly) persist even if its object isn't worthy of it—i.e., even if there's nothing about our beloveds in virtue of which they *merit* the intense love that we feel for them. And I think that similarly plausible things can be said to mitigate any residual, apparent weirdness that might seem to attach to certain implications of the kind of quality view that I've been defending in this paper.

6.2. *Trading Up Redux*

I've argued that quality theorists can answer the problem of trading up in the same way that they can answer the rest of the "big four" problems standardly raised against them—by rejecting an equivalence between the fittingness relation and the reason relation, and by supplementing their account of what makes love fitting with the substantive normative claim that valuable, loving relationships provide (derivative) value-related reasons for love. In a situation in which, intuitively, I shouldn't trade up, what explains this is the presence of value-related reasons, which are provided, ultimately, by the value of my current relationship. But suppose that someone is more lovable than my beloved and my relationship with that person would be even *better* than my current one. Why shouldn't I trade up then? This is the problem of trading up redux.

I'll offer two distinct and incompatible types of response to this problem. I think each type of response provides an adequate answer to the objection, but I want to stay neutral as to which is overall preferable. I leave it to the reader to choose her own adventure.

The first type of response to the problem of trading up redux is what I call the "diehard romantic" response. According to the diehard romantic, the problem of trading up redux falsely presupposes that a relationship with someone else *could* be better than the relationship you currently have with your beloved. This is because, according to the diehard romantic, valuable loving relationships have a kind of value that defies comparison—in Kant's terminology, they possess a *dignity*, as opposed to a *price*.[32] And

[32] Kant says that morality (and humanity inasmuch as it's capable of morality) is the only thing that has (or could possibly have) a dignity rather than a price (1785/1964, 4:435). But I see no reason why other kinds of entity couldn't also possess this sort of value; and Velleman (1999) and Anderson (1993) agree.

as a consequence, the kind of comparison the possibility of which the problem of trading up redux presupposes is in fact impossible—the value of your loving relationship is incomparable with the actual or prospective value of any other. Or to put the idea more precisely: there exists no actual or possible valuable loving relationship that bears any relation of comparability to the valuable loving relationship that you're currently in—your valuable loving relationship is neither worse, nor better, than any actual or possible alternative, and no other actual or possible alternative is, or could be, equally good. And what this means (*inter alia*) is that value-related reasons provided by a prospectively valuable relationship can't possibly favor loving someone more than, or instead of, someone with whom you already have a valuable relationship.[33] And so for the diehard romantic, the problem of trading up redux in fact presents no problem at all.

The diehard romantic response is not implausible. We certainly tend to treat and regard our most valued loving relationships as incomparable with potential alternatives. Moreover, we tend to value our loving relationships (and the loving relationships of others) in precisely the ways that Kant (and various Kantians) think that it's fitting for us to value things with a dignity—we respect, esteem, and revere them.[34] So perhaps the correct response to the problem of trading up redux is that of the diehard romantic.

Still, some will likely find the diehard romantic response objectionably saccharine. Is there really no possible situation in which we might permissibly pursue someone who might be a "better match?" Or some might simply prefer a less axiologically extravagant proposal—one which doesn't posit the special, incomparable kind of value that Kant calls a "dignity." So to those who might harbor such worries, let

---

[33] I'm here assuming (plausibly, I think) that the weights of value-related reasons provided by valuable relationships correspond to the degree to which the relationships that provide the reasons are valuable.
[34] See Kant (1785/1964, 4:428), and, in particular, Velleman (1999) and Anderson (1993, chap. 1).

me recommend an alternative response to the problem of trading up redux: what I call the "austere realist" response.

In contrast to the diehard romantic, the austere realist simply accepts the apparently problematic implication to which the problem of trading up redux points: if a new loving relationship with someone else would be discernably better than the loving relationship that you currently have with your beloved, then, other things equal, you have decisive all-things-considered reason to trade up. (Alternatively if a new loving relationship would be *equally* valuable to your current one, then the austere realist accepts that, other things equal, you have sufficient all-things-considered reason to trade "sideways.")

Now to the diehard romantic, the austere realist response will likely seem harsh. But there are things that can be said to mitigate its initial abrasiveness. First, other things are not always equal. In many cases, ending a valuable relationship with your beloved in favor of a new relationship with someone else will come with heavy costs. The pains of heartbreak and loss are real. And at least sometimes, the great disvalue associated with these pains (and other possible costs of ending a valuable loving relationship) will be enough to tip the scales in favor of "making it work" with your current beloved.[35]

Second, while the austere realist accepts that in some cases the value of your current relationship can be compared against the value of a possible new one, she needn't accept that this is so in all cases. For the kinds of considerations in virtue of which a loving relationship might qualify as valuable are variegated. And in cases in which the

[35] This kind of calculation might seem cold, but I think it's legitimate. This is why we don't tell people we no longer love them on their birthdays, or right after they lose a loved one. It also makes sense of the phenomenon in which parents stay together, and try to continue loving each other, "for the sake of the kids."

good-making features of your current relationship differ considerably from the good-making features of your prospective relationship, making overall, all-things-considered comparisons between the two relationships might be impossible. For example, your current relationship might be a better relationship in the respect that it fosters mutual encouragement between you and your beloved to pursue the plans and projects to which you're dedicated. And your prospective relationship might be a better relationship in the respect that it would facilitate your developing a more carefree, relaxed attitude toward life's pressures. So your current relationship is in some respects a better relationship than your prospective one would be, and your prospective relationship would be in some respects a better relationship than your current one is. But it might well be that there's simply no truth about how the two relationships compare *tout court*—that your current relationship is not better, worse, or equally good, as compared to your prospective one. And if this is so, then the value-related reasons provided by your prospective relationship can't favor loving your potential beloved more than, or instead of, your current beloved. So the austere realist can say that in a large class of cases—viz., cases in which the loving relationships in play qualify as valuable in virtue of rather different kinds of considerations—her proposal won't recommend (or permit) that you trade up.

So the austere realist response isn't as harsh as it may initially seem. Still, it does predict that there are possible cases in which you have decisive (or sufficient) reason to trade up. But maybe this is exactly right. After all, the austere realist response predicts that you ought to trade up only if a new relationship would be discernably better than your current one, and the value of the new relationship would be at least sufficient to outweigh whatever disvalue might be associated with leaving your current relationship.

Now this may be unromantic, but it doesn't sound all that implausible. Perhaps, contrary to the diehard romantic, if it truly would be better on all counts to be with someone other than your current beloved, then you really ought, all-things-considered, to be with him.

Now as I said, I want to remain officially neutral between the austere realist response and that of the diehard romantic. The right route isn't clear to me. But as I also said, I think that either type of response is sufficient to answer the problem of trading up redux. So, I don't think this problem speaks decisively against the proposal of this paper.

REFERENCES

Abramson, Kate and Adam Leite. “Love as a Reactive Emotion.” *Philosophical Quarterly* 61 (2011): 673-99.

Alston, William. *Beyond Justification: Dimensions of Epistemic Evaluation*. Ithaca, NY: Cornell University Press, 2005.

Anderson, Elizabeth. *Value in Ethics and Economics*. Cambridge, MA: Harvard University Press, 1993.

Badhwar, Neera Kapur. “Friends as Ends in Themselves.” *Philosophy and Phenomenological Research* 48 (1987): 1-23.

Bagley, Benjamin. “Loving Someone in Particular.” *Ethics* (2015): 477-507.

Baker, Derek. “Skepticism about Ought *Simpliciter*.” Unpublished manuscript.

Beardsley, Monroe C. “Intrinsic Value.” *Philosophy and Phenomenological Research* 26 (1965): 1-17.

Berker, Selim. “A Combinatorial Argument against Practical Reasons for Belief.” Unpublished manuscript.

Brandt, Richard. “Moral Valuation.” *Ethics* 56 (1946): 106-21.

Brentano, Franz. *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, John. *Rationality Through Reasoning* Hoboken, NJ: Wiley-Blackwell, 2013.

---. “Reasons.” In *Reason and Value: Themes from the Moral Philosophy of Joseph Raz*, edited by Jay Wallace, Michael Smith, Samuel Scheffler, and Philip Pettit, 228-55. Oxford: Oxford University Press, 2004.

---. “Reply to Southwood, Kearns and Star, and Cullity.” *Ethics* 119 (2008): 96-108.

Brown, Robert. *Analyzing Love*. Cambridge: Cambridge University Press, 1987.

Brunero, John. “Reasons, Evidence, and Explanations.” In *The Oxford Handbook of Reasons and Normativity*, edited by Daniel Star. Oxford: Oxford University Press, forthcoming.

Chang, Ruth. "Introduction." In *Incommensurability, Incomparability and Practical Reason*, edited by Ruth Chang, 1-34. Cambridge, MA: Harvard University Press, 1997.

Chappell, Richard. "Fittingness: The Sole Normative Primitive." *Philosophical Quarterly* 62 (2012): 684-704.

Chisholm, Roderick. *Brentano and Intrinsic Value*. Cambridge: Cambridge University Press, 1986.

Clarke, Randolph, Michael McKenna, and Angela M. Smith, eds. *The Nature of Moral Responsibility: New Essays*. Oxford: Oxford University Press, 2015.

Clarke, Randolph. "Some Theses on Desert." *Philosophical Explorations* 16 (2013): 153-64.

Clausen, Ginger. "Love and Organic Unities." Ph.D. Diss., University of Arizona, 2016.

Cohen, Stewart. "Reasons to Believe and Reasons to Act." *Episteme* 13 (2016) 427-38.

Copp, David. "The Ring of Gyges: Overridingness and the Unity of Reason." *Social Philosophy and Policy* 14 (1997): 86-101.

Crisp, Roger. "Review of John Kupperman, Value…and What Follows." *Philosophy* 75 (2000): 458-62.

Dancy, Jonathan. "Should We Pass the Buck?" In *Philosophy, the Good, the True, and the Beautiful*, edited by Anthony O'Hear, 159-74. New York: Cambridge University Press, 2000.

Danielsson, Sven and Jonas Olson. "Brentano and the Buck-Passers." *Mind* 116 (2007): 511-22.

D'Arms, Justin and Daniel 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.

Delaney, Neil. "Romantic Love and Loving Commitment: Articulating a Modern Ideal." *Philosophical Quarterly* 33 (1996): 339-56.

Ewing, A.C. *The Definition of Good*. London: Routledge & Kegan Paul, 1948.

Feinberg, Joel. *Doing & Deserving: Essays in the Theory of Responsibility*. Princeton, NJ: Princeton University Press, 1970.

Feldman, Richard. “The Ethics of Belief.” *Philosophy and Phenomenological Research* 60 (2000): 667-95.

Fine, Kit. “Essence and Modality.” *Philosophical Perspectives* 8 (1994): 1-16.

Finlay, Stephen. “Explaining Reasons.” *Deutsche Jahrbuch für Philosophie* 4 (2012): 112-26.

---.“The Reasons that Matter.” *Australasian Journal of Philosophy* 84 (2006): 1-20.

Foley, Richard. 1993. *Working Without a Net: A Study of Egocentric Epistemology*. New York: Oxford University Press.

---. “Justified Belief as Responsible Belief.” In *Contemporary Debates in Epistemology*, edited by Matthias Steup and Ernest Sosa, 313-26. Malden, MA: Blackwell Publishing, 2003.

Frankfurt, Harry G. *Necessity, Volition, and Love*. New York: Cambridge University Press, 1999.

---. *Taking Ourselves Seriously and Getting it Right*. Stanford, CA: Stanford University Press, 2006.

Friedman, Jane. “Suspended Judgment.” *Philosophical Studies* 162 (2013): 165-81.

Gibbard, Allan. *Wise Choices*, *Apt Feelings*. Cambridge, MA: Harvard University Press, 1990.

Goldman, Alvin. *Knowledge in a Social World*. Oxford: Oxford University Press, 1999.

Harman, Gilbert. “Practical Aspects of Theoretical Reasoning.” In *The Oxford Handbook of Rationality*, edited by Alfred R. Mele and Piers Rawling, 45-56. Oxford: Oxford University Press, 2004.

Heathwood, Christopher. “Fitting Attitudes and Welfare.” In *Oxford Studies in Metaethics*, vol. 3, edited by Russ Shafer-Landau, 47-73. Oxford: Oxford University Press, 2008.

Howard, Christopher M. “The Fundamentality of Fit.” In *Oxford Studies in Metaethics*, vol. 14, edited by Russ Shafer-Landau. Oxford University Press: forthcoming.

---. “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 (2016b): 1191-1201.

Hurka, Thomas. “The Justification of National Partiality.” In *The Morality of Nationalism,* edited by Robert McKim and Jeff McMahan, 139-57. New York: Oxford University Press, 1997.

---. *Virtue*, *Vice*, *and Value*. Oxford: Oxford University Press, 2001.

Hieronymi, Pamela. “The Wrong Kind of Reason.” *Journal of Philosophy* 102 (2005): 437-57.

Jackson, Frank. *From Metaphysics to Ethics*. Oxford: Oxford University Press, 1997.

James, William. *The Will to Believe and Other Essays in Popular Philosophy*. Cambridge, MA: Harvard University Press, 1979. Originally published in 1897.

Jollimore, Troy. *Love’s Vision*. Princeton, NJ: Princeton University Press, 2011.

Kagan, Shelly. *Normative Ethics*. Boulder, CO: Westview Press, 1998.

---. “Rethinking Intrinsic Value.” *Journal of Ethics* 2 (1998): 277-97.

Kant, Immanuel. *Groundwork of the Metaphysics of Morals*. Translated by H.J. Paton. New York: Harper & Row, 1964. Originally published in 1785.

Kavka, Gregory. “The Toxin Puzzle.” *Analysis* 43 (1983): 33-36.

Kearns, Stephen and Daniel Star. “Reasons: Explanations or Evidence?” *Ethics* 119 (2008): 31-56.

---.“Reasons as Evidence” In *Oxford Studies in Metaethics*, vol. 4, edited by Russ Shafer-Landau, 215-42. Oxford: Oxford University Press, 2009

Keller, Simon. “How do I love thee? Let me count the properties.” *Philosophical Quarterly* 37 (2000): 163-73.

---. *Partiality.* Princeton, NJ: Princeton University Press, 2013.

Kelly, Thomas. “The Rationality of Belief and Some Other Propositional Attitudes.” *Philosophical Studies* 110 (2002): 163-96.

Kolodny, Niko. “Love as Valuing a Relationship.” *Philosophical Review* 112 (2003): 135-89.

---. “Why Be Rational?” *Mind* 114 (2005): 509-63.

Korsgaard, Christine. “Two Distinctions in Goodness.” *Philosophical Review* 92 (1983): 169-95.

Leary, Stephanie. “In Defense of Practical Reasons for Belief.” *Australasian Journal of Philosophy*, forthcoming.

Lewis, David. “Extrinsic Properties.” *Philosophical Studies* 44 (1983): 197-200.

Louise, Jennie. “Correct Responses and the Priority of the Normative.” *Ethical Theory and Moral Practice* 12 (2009): 345-64.

Lynch, Michael. *True to Life: Why Truth Matters*. Cambridge, MA: Massachusetts Institute of Technology Press, 2004.

Maguire, Barry. “There are no Reasons for Emotions or Desires.” Unpublished manuscript.

---. “The Value-Based Theory of Reasons.” *Ergo* 3 (2016): 233-62.

Markovits, Julia. “Internal Reasons and the Motivating Intuition.” In *New Waves in Metaethics*, edited by Michael Brady, 141-165. Basingstoke: Palgrave Macmillan, 2011.

---. *Moral Reason.* New York: Oxford University Press, 2014.

Marušić, Berislav. *Evidence and Agency: Norms of Belief for Promising and Resolving*. Oxford: Oxford University Press, 2016.

McDowell, John. “Values and Secondary Qualities.” Reprinted in *Mind, Value and Reality*. Cambridge, MA: Harvard University Press, 1998.

McHugh, Connor and Jonathan Way. “Fittingness First.” *Ethics* 126 (2016): 575-606.

McKeever, Sean and Michael Ridge. “Elusive Reasons.” In *Oxford Studies in Metaethics*, vol. 7, edited by Russ Shafer-Landau, 110-37. Oxford: Oxford University Press, 2012.

McLeod, Owen. “Just Plain ‘Ought’.” *Journal of Ethics* 5 (2001): 269-91.

Moore, G.E. *Principia Ethica*. Cambridge: Cambridge University Press, 1903.

Naar, Hichem. “Subject-Relative Reasons for Love.” *Ratio* 30 (2017): 197-214.

Nair, Shyam. “How Do Reasons Accrue?” In *Weighing Reasons*, edited by Errol Lord and Barry Maguire, 56-73. Oxford: Oxford University Press, 2016.

Nozick, Robert. “Love’s Bond.” In *The Examined Life: Philosophical Meditations*, 68-86. New York: Simon & Schuster, 1989.

Owens, David. *Reasons without Freedom*. Oxford: Routledge, 2000.

Papineau, David. “There are No Norms of Belief.” In *The Aim of Belief*, edited by Timothy Chan, 64-79. Oxford: Oxford University Press, 2013.

Parfit, Derek. *On What Matters*. Oxford: Oxford University Press, 2011.

Pascal, Blaise. *Pensées*. Translated by W.F. Trotter. London: Dent, 1910. Originally published in 1670.

Piller, Christian. “Content-Related and Attitude-Related Reasons for Preferences.” *Royal Institute of Philosophy Supplement* 81 (2006): 155-81.

Portmore, Doug. “The Teleological Conception of Practical Reasons.” *Mind* 120 (2011): 117-53.

Protasi, Sara. “Loving People for Who They Are (Even When They Don’t Love You Back).” *Journal of Philosophy* 24 (2016): 214-34.

Rabinowicz, Wlodek and Toni Rønnow-Rasmussen. “The Strike of the Demon: On Fitting Pro-Attitudes and Value.” *Ethics* 114 (2004): 391-423.

Rawls, John. *A Theory of Justice*. Cambrdige, MA: Belknap Press, 1971.

Raz, Joseph. *Engaging Reason*. Oxford: Oxford University Press, 1999.

---“The Myth of Instrumental Rationality.” *Journal of Ethics and Social Philosophy* 1 (2005): 1-28.

---. *From Normativity to Responsibility*. Oxford: Oxford University Press, 2011.

---. “Reasons: Practical and Adaptive.” In *Reasons for Action*, edited by David Sobel and Steven Wall, 37-57. Cambridge: Cambridge University Press, 2009.

Reisner, Andrew. “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.

Rinard, Susanna. “Against the New Evidentialists.” *Philosophical Issues* 25 (2015): 208-23.

---. “No Exception for Belief.” *Philosophy and Phenomenological Research* 94 (2017): 121-43.

Rosen, Gideon. "The Alethic Conception of Moral Responsibility." In *The Nature of Moral Responsibility: New Essays*, edited by Randolph Clarke, Michael McKenna, and Angela 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.M. *What We Owe to Each Other*. Cambridge, MA: Harvard University Press, 1998.

Schroeder, Mark. *Slaves of the Passions*. Oxford: Oxford University Press, 2007.

---. "Teleology, Agent-Relative Value, and 'Good'." *Ethics* 117 (2007): 265-300.

---. "Value and the Right Kind of Reason." In *Oxford Studies in Metaethics*, vol. 5, edited by Russ Shafer-Landau, 25-55. Oxford: Oxford University Press, 2010.

---. "The Ubiquity of State-Given Reasons." *Ethics* 122 (2012): 457-88.

---. "What Makes Reasons Sufficient?" *American Philosophical Quarterly* 52 (2015): 159-70.

Setiya, Kieran. "Love and the Value of a Life." *Philosophical Review* 123 (2014): 251-80.

Shah, Nishi. "A New Argument for Evidentialism." *Philosophical Quarterly* 56 (2006): 481-98.

---. "How Action Governs Intention." *Philosopher's Imprint* 8 (2008): 1-19.

Sharadin, Nathaniel. "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, John. *The Domain of Reasons*. Oxford: Oxford University Press, 2010.

Smith, Michael. "Two Kinds of Consequentialism." *Philosophical Issues* 19 (2009): 257-72.

Smuts, Aaron. "In Defense of the No-Reasons View of Love." Unpublished manuscript.

Snedegar, Justin. “Negative Reason Existentials.” *Thought* 2 (2013): 108-116.

Stang, Nicholas. “Artworks Are Not Valuable for Their Own Sake.” *Journal of Aesthetics & Art Criticism* 70 (2012): 271-80.

Steglich-Peterson, Asbjørn. “How to be a Teleologist about Epistemic Reasons.” In *Reasons for Belief*, edited by Asbjørn Steglich-Peterson and Andrew Reisner, 13-33. Cambridge: Cambridge University Press, 2011.

Stratton-Lake, Philip. “How to Deal with Evil Demons: Comment on Rabinowicz and Rønnow-Rasmussen.” *Ethics* 115 (2005): 788-98.

Suikkanen, Jussi. “Consequentialism, Constraints, and the Good-Relative-To: A Reply to Mark Schroeder.” *Journal of Ethics and Social Philosophy* (March 2009): 1-9.

Svavarsdottir, Sigrun. “Having Value and Being Worth Valuing.” *Journal of Philosophy* 111 (2014): 84-109.

Sylvan, Kurt. “Truth Monism Without Teleology.” *Thought: A Journal of Philosophy* 1 (2012): 161-63.

Thomson, Judith Jarvis. *Normativity*. Chicago and La Salle, IL: Open Court, 2008.

Tiffany, Evan. “Deflationary Normative Pluralism.” *Canadian Journal of Philosophy* 37 (2007): 231-62.

Tucker, Miles. “The Pen, the Dress and the Coat: A Confusion in Goodness.” *Philosophical Studies* 173 (2016): 1911-22.

Velleman, J. David. “Love as a Moral Emotion.” *Ethics* 109 (1999): 338-74.

Way, Jonathan. “Transmission and the Wrong Kind of Reason.” *Ethics* 122 (2012): 489-515.

---. “Two Arguments for Evidentialism.” *Philosophical Quarterly* 66: 805-18.

---. “Value and Reasons to Favour.” In *Oxford Studies in Metaethics*, vol. 8, edited by Russ Shafer-Landau, 27-49. Oxford: Oxford University Press, 2013.

Wedgwood, Ralph. *The Nature of Normativity*. Oxford: Oxford University Press, 2007.

Whiting, Daniel. “Reasons for Belief, Reasons for Action, the Aim of Belief, the Aim of Action.” In *Epistemic Norms: New Essays on Action, Belief, and Assertion*, edited by Clayton Littlejohn and John Turri, 219-38. New York: Oxford University Press, 2014

Wiggins, David. “A Sensible Subjectivism?” In *Needs, Values, Truth: Essays in the Philosophy of Value*. Oxford: Blackwell, 1987.

Williams, Bernard. “Internal and External Reasons.” Reprinted in *Moral Luck*. Cambridge: Cambridge University Press, 1981.

Zangwill, Nick. “Love: Gloriously Amoral and Arational.” *Philosophical Explorations* 16 (2013): 298-314.

Zimmerman, Michael. “Partiality and Intrinsic Value.” *Mind* 120 (2011): 447-83.

---. “Varieties of Moral Responsibility.” In *The Nature of Moral Responsibility: New Essays*, edited by Randolph Clarke, Michael McKenna, and Angela M. Smith, 45-65. Oxford: Oxford University Press, 2015.