

## **Ambidextrous Reasons (or Why Reasons First's Reasons Aren't Facts)**

**Abstract:** The wrong kind of reason (WKR) problem is a problem for attempts to analyze normative properties using only facts about the balance of normative reasons, a style of analysis on which the 'Reasons First' programme depends. I argue that this problem cannot be solved if *the orthodox view of reasons* is true --- that is, if each normative reason is numerically identical with some fact, proposition, or state-of-affairs. That's because solving the WKR problem requires completely distinguishing between the right- and wrong-kind reasons for an attitude. I argue that some facts give both right- and wrong-kind reasons for an attitude. Consequently, no such distinction between the two types of reasons is complete if reasons are facts or the like. I conclude by suggesting that reasons and facts are related by constitution, not identity.

**Keywords:** Wrong Kind of Reason Problem, Weighing explanations, Buck-passing, Reasons First

# Ambidextrous Reasons (or Why Reasons First's Reasons Aren't Facts)

## 1. Introduction

When deliberating about what to do, we sometimes weigh an action's pros against its cons. Doing so resembles a way of analyzing normative properties, such as an act's rightness or a belief's justification, known as *weighing explanation*.<sup>1</sup> For example, we might hold that something is good, very roughly, just when and because the reasons given by its good-making features outweigh or beat the reasons given by its countervailing bad-making features. While appealing, the main problem with weighing explanations is distinguishing the reasons to be weighed from those to be ignored without presupposing claims about the normative feature to be analyzed. For example, we cannot identify the reasons that belong in the weighing analysis of goodness as the ones that make something good, as I've just done, on pain of triviality. The problem of providing a non-circular rationale for excluding certain reasons from analysis is what's known as the *wrong kind of reason* (WKR) problem.

This paper argues that this problem cannot be solved if each normative reason is numerically identical to some fact, proposition, or state-of-affairs.<sup>2</sup> I'll call this position *the orthodox view of reasons*. When the normative reason to donate to Oxfam is that doing so will save a life, the orthodox view holds that the reason to donate to Oxfam is numerically identical to the fact that doing so will save a life. Since weighing explanations succeed only if the WKR problem can be solved, this paper argues that weighing explanations and the orthodox view are incompatible.

The thesis's narrow phrasing belies its broad importance. Many philosophers are particularly taken with weighing explanations. According to these 'Reasons Firsters' all (or nearly all) normative properties have a weighing explanation. As such, the WKR problem is an especially serious challenge to Reasons First. However, virtually all Reasons Firsters assume the orthodox view of reasons. For example, T.M. Scanlon, Derek Parfit, Jonathan Dancy, and (pre-2018) Mark Schroeder all combine Reasons First with the orthodox view. As a result, the argument I offer below shows that a (if not *the*)

---

<sup>1</sup> Although W. D. Ross's *The Right and the Good* offers, arguably, the first discussion of weighing explanation, contemporary discussion of such explanations owes much to John Broome --- see Broome (2013) and Kearns and Star (2015).

<sup>2</sup> For brevity, I will refer to the orthodox view throughout as the view that reasons are facts, but this use is elliptical for the more cumbersome but accurate claim that the orthodox view is the view that reasons are facts, propositions, or states of affairs. Here's a representative sample of the orthodoxy: "Reasons are facts in that they are true propositions: they are the sort of things that can be known or believed, and by which it is rational to be moved" (Setiya (2015: 450)); "Objective reasons are facts which count in favour of (or against) some response" (McHugh (2018: 160)); "For all propositions  $r$ , agents  $x$ , and actions  $a$ , if  $r$  is a reason for  $x$  to do  $a$ , that is because there is some  $p$ ..." (Schroeder (2007: 29)); "... "is a reason for" is a four-place relation,  $R(p, x, c, a)$ , holding between a fact  $p$ , an agent  $x$ , a set of conditions  $c$ , and an action or attitude  $a$ . This is the relation that holds just in case  $p$  is a reason for a person  $x$  in situation  $c$  to do or hold  $a$ " (Scanlon (2014: 32)). It has been pointed out to me that the claim 'reasons are facts' is ambiguous between the idea that reasons are *identical* to facts and that reasons are *constituted* by facts (*c.f.*, 'those statues are clay; those are marble'). [Redacted], [redacted], and [redacted] confirm in personal correspondence that they intend the orthodox reading of that claim. Heterodox views exist. For example, Crisp (2006: 61-2) holds that reasons are properties of actions and Nagel (1970: 47) holds that they are predicates. I bracket such views for the remainder of the paper.

dominant position in contemporary metaethics is incoherent. Indeed, many challenges thought to afflict Reasons First actually originate in the incompatibility of Reasons First and the orthodox view.<sup>3</sup>

I develop this argument over five sections. Sections two and three show that solving the WKR problem requires articulating a description that not only *includes* all of the reasons that are right for analyzing a given normative property ('right-kind reasons') but also *excludes* all of the reasons that are wrong for that purpose ('wrong-kind reasons'). Section four shows that doing so is logically impossible if the orthodox view is true.

The explanation of why it's impossible is surprisingly simple. An overlooked kind of fact, 'ambidextrous facts' as I call them, gives both right- and wrong-kind reasons relative to an analysis. The orthodox view implies that these reasons are of both the right and wrong kinds, simultaneously. Consequently, if we assume the orthodox view, solving the WKR problem requires simultaneously including *and* excluding these ambidextrous facts from analysis, which is impossible, just as you cannot simultaneously invite and not invite someone to a party or let the dog in while letting him out. Section five rebuts two responses to the argument.

Ambidextrous facts, therefore, force a choice between the idea that reasons are numerically identical to facts and the style of analysis employed by Reasons First. Philosophers already sceptical of Reasons First may welcome this conclusion as yet more evidence for their scepticism. Doing so is consistent with what follows. However, section six concludes by sketching a heterodox view of the relationship between facts and reasons that is compatible with Reasons First, recommending it to those who wish to pursue research in that programme.

## 2. Setup

Reasons First holds that only certain facts --- namely, facts about the competition between normative reasons --- ultimately explain normative properties; reasons are thus 'first' in the sense of metaphysically preceding all other normative properties. Reasons Firsters often express this competition through the metaphor of weight. For proponents of Reasons First, reasons' weighing role is distinctive:

*Reason:* An object *o* is a normative reason iff *o* is the kind of thing weighed in a weighing explanation of a normative property.

As I intend it, *Reason* is nothing more than a helpful truism.<sup>4</sup> It fixes discussion by describing normative reasons functionally, in terms of what they *do*, not in terms of what they *are*. This functional characterization will help us to assess which things can be reasons --- the question at this

---

<sup>3</sup> For example, I argue in Howard (fc-a) that a well-known family of challenges to a reasons-based analysis of moral worth are actually challenges to the orthodox view.

<sup>4</sup> In particular, I do not advance *Reason* as an analysis. Doing so needlessly courts the controversy of whether reasons are primitive, unanalyzable entities. I take no stand on that issue here.

paper's heart --- by assessing whether those things can play this weighing role. As I'll argue, it turns out that, surprisingly, neither facts nor propositions nor states of affairs are fit to play this role, so they cannot be Reasons First's reasons.

A paradigm example of a weighing analysis is Scanlon (1998)'s well-known "buck-passing" analysis of final (or non-instrumental) goodness. Here is one way of articulating this style of analysis:<sup>5</sup>

*Buck-Passing Analysis\** (BPA\*): An object is *finally good* just when and because there is sufficient reason to value it for its own sake.<sup>6</sup>

'Sufficient' is a term of art. A set of reasons is sufficient, relative to a set of countervailing reasons, just when the first set is at least as strong or weighty as the second. As a result, BPA\* holds that an object is finally good just when and because the reasons for valuing it for its own sake are at least as weighty the reasons against valuing it for its own sake. This makes BPA\* a paradigmatic weighing explanation. Analyses like BPA\* do not entail Reasons First since they leave open whether their right-hand sides admit of further analysis. But Reasons First entails analyses like these, so trouble for the analyses is trouble for Reasons First.

Indeed, BPA\* isn't quite right. For example, Crisp (2000), echoing Anscombe, asks us to imagine a demon who threatens to inflict severe pain on you unless you value an otherwise worthless saucer of mud for its own sake. The threat gives a very strong reason to value it.<sup>7</sup> Moreover, if the relevant countervailing reasons against doing so are not very strong, then there is sufficient reason to value the saucer of mud for its own sake. But clearly the demon's threat doesn't make the saucer finally good. Cases like these falsify the sufficiency direction of BPA\*'s analysis.

As a result, BPA\*'s right-hand side must be strengthened to be sufficient. In particular, it must be strengthened to include all the reasons for valuing something that contribute to its final goodness and to exclude all the reasons that do not:

*Buck-Passing Analysis* (BPA): An object is *finally good* just when and because there is a sufficient reason of the right kind to value it for its own sake.

While the demon's threat offers a sufficient reason to value the saucer, the threat doesn't contribute to the saucer's final goodness, so it does not offer the right kind of reason for the BPA. Reasons given by

---

<sup>5</sup> There are several ways of understanding Scanlon's project, as Liao (2010) canvasses. I favour the biconditional understanding of the BPA --- as does Dancy (2000: 168) --- as it foregrounds my criticism of the orthodox account of reasons.

<sup>6</sup> As Dancy (2000: 168) and Crisp (2005: 82) argue, identifying precisely which attitude to associate with the BPA is a challenge. I'm using 'value' stipulatively to name that attitude, whatever it is. See also Gert (2016).

<sup>7</sup> Some, such as Hieronymi (2005) and Way (2012), deny that there is a WKR problem by denying that the demon's threat is a reason to value the saucer. These WKR skeptics argue that the threat is a reason to bring it about that one so values the saucer rather than a reason to value it. For some rebuttals, consider Reisner (2009), Howard (2016), and Leary (2018). I'll bracket WKR skepticism for the moment, returning to it in footnote seventeen.

the threat must be excluded from the set of reasons whose sufficiency necessitates final goodness, known as the ‘right reasons’ for the BPA.

However, while it is intuitive that the demon’s threat does not make the saucer finally good, we must appeal to more than intuition to explain why the threat doesn’t give right-kind reasons to value the saucer. Moreover, in doing so, we cannot appeal to claims about final goodness, as I’ve just done in immediately above. After all, weighing analyses like the BPA use claims about right-kind reasons to analyze final value. We cannot then analyze right-kind reasons using final value without trivializing both analyses. The challenge of non-circularly excluding normative reasons like those given by the threat is known as the WKR problem. It is the central problem for weighing explanations.

The WKR problem plausibly afflicts more than just the BPA. On one understanding, the BPA tacitly conjoins two claims: that something is finally good just when it is *fitting* to value it for its own sake, and that valuing something for its own sake is fitting just when there is sufficient (right-kind) reason to value it in that way. The demon’s threat challenges the second connection, providing sufficient reason to value something that is not fitting to value.

Connecting reasons and fittingness in this way expands the WKR problem’s scope beyond the BPA. For instance, we might hold that a belief is fitting just when it is *propositionally justified*,<sup>8</sup> and hold that a belief is propositionally justified just when and because there is sufficient evidential reason for it:

*Propositional Justification* (PJ): A belief is *propositionally justified* just when and because there is a sufficient evidential reason to have it.

‘Evidential’ plays the same role in PJ that ‘of the right kind’ plays in the BPA, restricting the reasons used by the analysis to only those that bear on the normative property to be analyzed. However, just as we must understand ‘right-kind’ in the context of the BPA independently from claims about final value, we must understand ‘evidence’ in the context of PJ independently from claims about propositional justification. Without such a restriction, PJ is vulnerable to counterexamples structurally analogous to ones that afflict the BPA. For example, if the demon threatens to destroy the world unless you believe that the moon is made of green cheese, you have a very strong reason for a belief that is not justified. So we must restrict PJ’s right-hand side to only evidential reasons for belief and give a characterization of evidence that does not involve propositional justification.

But we may interpret the WKR problem even more broadly still, beyond analyses of fitting attitudes.<sup>9</sup> We can hold that any analysis of a ‘standard of correctness’, such as being a morally required action,

---

<sup>8</sup> Of course, there are other ways for a belief to be fitting: when it is true or when it is knowledge. These other kinds of fitting belief are not central here, so I will bracket them. In my stipulative use ‘fitting belief’ means ‘propositionally justified belief’.

<sup>9</sup> Some philosophers’ understanding of the WKR problem, like Heuer’s (2010), restricts it to fitting attitudes accounts. I interpret the problem more broadly to include analyses like MR. Not much turns on this choice. The problem I raise arises on both the narrow and broad conceptions.

that depends on a restriction to certain reasons is susceptible to a WKR problem.<sup>10</sup> For example, consider the following:

*Moral Requirement (MR):* An act is *morally required* just when and because there is a conclusive moral reason to do it.<sup>11</sup>

MR depends on restricting its reasons to moral reasons. Relaxing this restriction invites counterexamples. For example, if going for a walk would please me, then there is a reason for me to go for the walk. But it is not ordinarily thought that this fact enjoins a moral requirement to go walking. That's because the pleasure I derive from walking plausibly gives me a merely prudential, not a moral, reason to walk. MR's weighing analysis must exclude these merely prudential reasons.<sup>12</sup> But it must provide an independent criterion by which to exclude them, much as PJ must provide a criterion for being evidence that is independent from notions of propositional justification or as the BPA must provide an independent criterion for being the right kind of reason to value something.

### 3. Three Conditions on Solving the WKR Problem

We've just seen how most weighing analyses involve only certain normative reasons and not others. The WKR problem for these weighing analysis is the problem of specifying the reasons that matter to the analysis without undermining it. In particular, we cannot appeal to the normative property analyzed when specifying these reasons. Doing so trivializes the analysis through circularity. For example, we cannot claim that the right-kind reasons for the BPA are those that involve an object's goodness if we also want to analyze goodness with claims about those reasons.

Rather, we must provide a description of the right reasons for an analysis, *D*, that meets three conditions. I've just mentioned the first:

1. *Non-triviality:* *D* doesn't refer to the normative property to be analyzed, on pain of circularity.

Fortunately, there's no shortage of proposals that claim to satisfy *Non-triviality*. Some appeal to differences in what grounds a normative reason's normativity. For example, as Parfit (2001) and Piller (2006) observe, bribes are reasons for an attitude like, for example, believing *P* in virtue of facts about the mental *state* of believing *P*, not in virtue of facts about *P*. This metaphysical distinction permits us to argue that only the latter, 'object-given' reasons are right-kind reasons for the PJ. Others appeal to a difference in motivational profile. For example, Dancy (2000), Hieronymi (2005), Raz (2011), and Rowland (2015) argue that only reasons for which we can hold a belief are the right kind of reasons for that belief. Many wrong-kind reasons appear not to meet this condition --- it's hard to voluntarily believe *P* as a response to a threat or a bribe to believe *P*. Alternatively, Way (2012) traces differences in

---

<sup>10</sup> See, for example, Danielsson and Olson (2007), Schroeder (2010), and Sharadin (2013).

<sup>11</sup> By 'conclusive', relative to a set of countervailing reasons, I simply mean being more weighty than the countervailing set.

<sup>12</sup> If we accept the possibility of supererogation, the interaction between prudential and moral reasons is more complicated than I've suggested here. I'm bracketing that possibility to simplify discussion.

how certain reasons “transmit” by making other facts into reasons, which map onto intuitive differences between right- and wrong-kind reasons.

These proposals also implicitly meet the second condition on articulating *D*. We could satisfy non-triviality *and* pick out the right reasons simply by offering a long conjunction like “*that* reason, and *that* reason, but not *that* reason or *that* reason...” But the Reasons First approach analyzes normative properties with the aim of understanding them. Picking out the right-kind reasons purely demonstratively doesn’t enlighten us about how the analysis works since it doesn’t explain why only those reasons are suited to it. As a result, proposals for *D* must meet a second condition:

2. *Non-vacuity*: *D* explains why the right reasons are the right ones for the analysis.

Finally, *D* must also meet a third condition, which is the most important one for our purposes here. The right-kind reasons for an analysis are the ones that must be included in it and wrong-kind reasons are the ones that must be excluded from it. If we fail to include all the right-kind reasons, we risk underweighting the set of right-kind reasons. For example, if some object is made finally good by only a single feature and we neglect the reasons given by that feature, we’ll wrongly judge that the object is not finally good for we’ll wrongly judge that there’s no reason to value it. Likewise, if we fail to exclude all the wrong-kind reasons, we risk overweighting the set of right-kind reasons. For example, as we’ve already seen, including the reasons given by threats and bribes can easily distort the balance of reasons. This gives us our third condition:

3. *Extensionality*: *D* includes all the reasons that are right for the analysis and excludes all the reasons that are wrong for it.

*Extensionality* is central to this paper’s argument. In the next section, I’ll show that *Extensionality* cannot be satisfied if the orthodox view is true, namely, if each normative reason is identical to the fact or proposition or state of affairs that gives it. This implication is particularly important because Reasons First requires solving the WKR problem, which requires satisfying *Extensionality*. Therefore, if the orthodox view makes *Extensionality* unsatisfiable, then it contradicts Reasons First.

## **4. Ambidextrous Facts and the Possibility of Ambidextrous Reasons**

### **4.1 Stating the Problem**

Consider an analogy. We cannot throw a party that both includes all football players and excludes all baseball players. Given that Bo Jackson was both a football *and* baseball all-star, Bo must be both included in the party, for he’s a football player, and excluded from it, for he’s a baseball player. But including and excluding Bo from the party at the very same time is impossible. So we cannot throw the party or succeed at any task that requires it.

*Extensionality* requires dividing reasons similarly to how the party requires dividing athletes, asking us to include some and to exclude others. As I'll now show, certain facts are like Bo Jackson in that they must be both included and excluded if they are reasons, making it impossible to satisfy *Extensionality*. For the rest of the paper, I'll refer simply to facts when discussing the orthodox view rather than the more cumbersome disjunction of facts or propositions or states of affairs. As I show in the conclusion, differences between these three kinds of entities are orthogonal to the basic problem that I raise.

The facts that make satisfying *Extensionality* incompatible with the orthodox view are what I call 'ambidextrous'. People are ambidextrous when they are adept with their right and left hands. Owing to the unfortunate but longstanding association between left-handedness and wrongness, a fact is *ambidextrous* in my usage when it gives both a right- and a wrong-kind reason, relative to an analysis.<sup>13</sup> For example, recall PJ:

*Propositional Justification* (PJ): A belief is *propositionally justified* just when and because there is a sufficient *evidential* reason to have it.

A right-kind reason for me to believe a claim, relative to PJ, is evidence that the claim is true. For example, when someone makes a generous charitable donation, that is evidence that they're rich --- other things equal, my confidence that someone is rich increases conditional on their having given a generous charitable donation. In contrast, a wrong-kind reason for me to believe something, relative to PJ, is an incentive to believe something, such as a bribe or threat. For example, if I'm offered a large sum of money to believe that someone is rich, then that bribe gives me a wrong-kind reason for belief.

But consider the following:

A. Bill will give you a large sum of money to believe that he's rich.

(A) is an ambidextrous fact. As such, it gives a right-kind reason in just the same way that a generous charitable donation does; it signals that the donor can afford to part with a large sum of money at no obvious comparable personal gain. Likewise, it gives a wrong-kind reason to believe that Bill is rich because it is an incentive to believe that he is rich, much as the demon's threat is an incentive to value the saucer of mud or to believe that the moon is made of green cheese. So (A) gives both a wrong- and a right-kind reason for you to believe that Bill is rich, relative to PJ.

Ambidextrous facts like (A) are entirely compatible *either* with weighing explanations like PJ *or* with the orthodox picture of reasons, if taken severally. However, ambidextrous facts show that the two views cannot be conjoined. If we assume the orthodox view, then both the right-kind and the wrong-kind reasons given by (A) are identical with it. And if each is identical with (A), then they are identical to each other by the transitivity of identity. So the orthodox view implies that (A) is a reason

---

<sup>13</sup> I am using 'give' stipulatively to name the relationship between facts and reasons, whatever it is. Advocates of the orthodox view hold that a fact gives a reason when it is identical to that reason. As we'll see in the concluding section, I take 'give' to mean something like 'partly constitute'.



of both the right and wrong kind relative to PJ, simultaneously. As a result, we cannot then exclude all the wrong-kind reasons from consideration in PJ while including all the right-kind reasons for doing so. This impossible requirement requires both including and excluding (A). Just as we cannot throw a party that excludes all baseball players and that includes all football players for doing so, it is impossible to invite and not invite Bo Jackson, we cannot satisfy *Extensionality* with respect to PJ for doing so. This impossible requirement requires both including and excluding (A) from PJ.

Ambidextrous facts highlight a counterintuitive consequence of the orthodox view. Intuitively, (A) gives *two*, numerically distinct reasons for you to believe that Bill is rich: one evidential and one wrong-kind or pragmatic. Intuitively, we can include the former when assessing PJ and exclude the latter. But (A) is only one fact. As a result, if the orthodox view is true, it gives only one reason for you to believe that Bill is rich. The orthodox view is thus counterintuitively stingy about how many reasons (A) gives for you to believe that Bill is rich.<sup>14</sup>

Observing this stinginess, philosophers sympathetic to the orthodox view may find it natural to say, “When I claim that reasons are facts, I don’t mean that reasons are *identical* to facts. I mean something else, like that facts and reasons are closely related.” This response is surely right. We *should* think that facts and reasons are closely related and we *shouldn’t* think that they’re related by identity. But those tempted by this response must recognize that it offers a rejection, not a defense, of the orthodox view, which holds that reasons and facts are identical. After all, according to proponents of the view, “reasons are facts in that they are true propositions: they are the sort of things that can be known or believed, and by which it is rational to be moved.”<sup>15</sup>

I’ve just argued that facts like (A) make satisfying *Extensionality* with respect to PJ impossible if we assume the orthodox view. My aim in the rest of this section is to expand this problem’s scope, showing that there’s a widespread structural incompatibility between the orthodox view and the full range of weighing analyses on which Reasons First depends. I rebut several responses to my argument in the next section and conclude by sketching an alternative to the orthodox view according to which facts and reasons are not related by identity but by constitution.

## 4.2 Assessing the Ubiquity of Ambidextrous Facts

Ambidextrous facts are common. Recall the BPA:

*Buck-Passing Analysis* (BPA): An object is *finally good* just when and because there is sufficient reason *of the right kind* to value it for its own sake.

As we’ve established, the right kind of reason for the BPA is one given by (final) good-making features. For example, that someone is benevolent is a good-making feature of their character. The wrong kind of reasons are given by incentives, like bribes or threats. For example, if I will save thousands of

---

<sup>14</sup> Some of these claims are anticipated by Maguire and Snedegar (forthcoming).

<sup>15</sup> Setiya (2015: 450)

children from starvation if you value a saucer of mud for its own sake, then that is a strong reason for you to value the saucer of mud for its own sake. Consequently, an ambidextrous fact relative to the BPA is one that combines a good-making feature of something with a threat or bribe to value it for its own sake.

Imagine the following scenario. Boris is a gifted public speaker and rhetorician. A group of Effective Altruists has invited Boris to speak on the topic of charitable donation, hoping he'll move the crowd to buy mosquito netting for sub-Saharan Africans. While giving the speech comes at a great cost to Boris --- he must pay his own way, miss his own birthday party, forgo great seats to the theatre, etc. --- those costs do not deter him because Boris cares deeply about preserving human life. However, Boris has succumbed to a bout of depression and self-doubt. He thinks himself worthless and cannot bring himself to speak in front of the crowd. Fortunately, Boris has great faith in your judgment. As a result, if you value him, Boris will give up his belief that he is worthless and give the speech. The only thing that stands between him and the speech is your opinion:

B. Boris will give the speech if you value him.

(B) gives a right-kind reason to value Boris for his own sake, relative to the BPA. That's because (B) expresses Boris's willingness to give the speech and, thus, his benevolent disposition. After all, the great costs of giving this speech, from which he will not personally profit, do not deter him from giving it. His virtuous desire to preserve human life drives him to give the speech, despite the sacrifices he must make to give it. As a result, (B) is a right-kind reason to value Boris.

However, (B) also gives a wrong-kind reason to value Boris for there is a strong incentive to do so: consider all of the sub-Saharan lives that will be saved as a result of Boris's considerable rhetorical talent. So (B) is an ambidextrous fact, relative to the BPA. Assuming the orthodox view then ensures that the BPA cannot satisfy *Extensionality* given (B).

Indeed, there's good evidence that *any* account of what makes an attitude fitting in terms of sufficient right-kind reasons for the attitude succumbs to ambidextrous facts, assuming the orthodox view. For example, some fears are fitting. Fear of global warming is fitting but fear of harmless puppies is not. Suppose that Ba'al-zebul, lesser cousin to Crisp's demon, is extremely insecure. He is jealous of the fear his cousin sows. To garner the fear he craves, Ba'al-zebul threatens to make full use of his power, inflicting an irritating and inconvenient series of traffic delays on Manhattan commuters. Hence:

C. Ba'al-zebul will impede Manhattan transit unless you fear him.

(C) is ambidextrous. It expresses both a fearsome display of demonic power and a powerful incentive to fear. Because it is ambidextrous, (C) undermines the claim that it is fitting to fear something just when and because there is sufficient right-kind reason to fear it. In Crisp's case, you have a strong reason to value something that is not valuable. Likewise, we can adjust the details in this case so that you have a strong reason to fear something that is not fearsome. Suppose that while Ba'al-zebul has

limited transit-inhibiting powers at his disposal, he is not all that fearsome. He is often overcome by his better nature, and secretly enjoys gardening and long walks on the beach. As a result, it is not fitting to fear Ba'al-zebul. Nevertheless, that doesn't imply that (C) isn't a very strong, right-kind reason to fear Ba'al-zebul. Suppose you live uptown and you must absolutely submit some immigration documents downtown today, otherwise you will be deported. In that case, you have a very strong reason to fear Ba'al-zebul if doing so is necessary for submitting the documents on time. Moreover, according to the orthodox view, (C) is a right-kind reason to fear Ba'al-zebul. So you have a very strong, right-kind reason to fear Ba'al-zebul, who is nevertheless not fearsome.

By way of closing out this section, observe that ambidextrous facts trouble even weighing analyses that do not involve fitting attitudes, either implicitly or explicitly, like the following:

*Moral Requirement (MR):* An act is *morally required* just in case and because there is a conclusive *moral* reason to do it.

The right-kind reasons for doing something, relative to MR, are given by facts that, to use W. D. Ross's expression, tend to make acting that way our duty proper or our moral duty. For example, that someone is in mortal need of aid is a moral reason to aid them. The wrong-kind reasons for doing something, relative to MR, are facts that tend to make it desirable to act some way independently from whether acting that way tends to be our moral duty. Facts from this familiar category involve incentives, like threats and bribes to act in certain ways. Consider the following claim:

D. Richie Rich is drowning in a lake.<sup>16</sup>

(D) gives a moral reason to jump in the lake, for it expresses the fact that Richie is in mortal need of aid. It also gives a non-moral reason to jump in the lake because you know that you'll get a fat reward for saving Richie. Thus (D) is an ambidextrous fact, which ensures that MR cannot satisfy *Extensionality* if the orthodox view is true.<sup>17</sup>

I've just argued that (A)-(D) are ambidextrous in the sense of giving right- and wrong-kind reasons, relative to various analyses. The reader may doubt this argument; indeed, I answer some such doubts in the next section. But it is important to note that my argument does not rest on showing that (A)-(D) are ambidextrous. It suffices that (A)-(D) make it *very likely* that *other* facts are ambidextrous, relative to the analyses above, even if (A)-(D) are not. Because the existence of ambidextrous facts is

---

<sup>16</sup> Howard (fc-b) provides an extended discussion of cases like these.

<sup>17</sup> Ambidextrous facts like (D) show that WKR scepticism, of the kind defended in Hieronymi (2005; 2013) and Way (2012), *inter alia*, is not a sufficiently general response to ambidextrous facts. Even in its non-moral capacity, (D) can clearly be a motivating reason for action; in both capacities, it bears on the question of what to do. So it cannot be excluded from consideration in MR on Hieronymi's reasoning. Likewise, even in its prudential capacity, (D) transmits in the right-kind way, so it is a right-kind reason for action by Way's lights too. As a result, WKR scepticism is not a sufficiently general response to the problem of ambidextrous reasons.

very likely, we should conclude that the orthodox view and the analyses above are incompatible, just as I've claimed, even in the face of quibbles about (A)-(D) themselves.

## 5. Objections and Responses

To organize various responses, I'll make my argument explicit. I assert:

1. The orthodox account is inconsistent with *Extensionality* for a wide range of analyses.
2. If an account of reasons cannot satisfy *Extensionality* for a wide range of analyses, then that account is incompatible with Reasons First.
3. Therefore, the orthodox account is incompatible with Reasons First.

Ambidextrous facts challenge *only* the combination of Reasons First and the orthodox view of reasons. If we deny that view and allow that the right- and wrong-kind reasons given by an ambidextrous fact are numerically distinct, ambidextrous facts are consistent with *Extensionality*. Likewise, if we deny Reasons First, we are free to analyze normative properties like propositional justification or value using claims that don't involve reasons. But it turns out that many important moral and normative theories combine Reasons First and the orthodox view and that certain objections that target Reasons First are actually objections to the combination of Reasons First and the orthodox view of reasons. That is why showing that they cannot be consistently combined is important.

### 5.1 Defending Premise (1)

#### 5.1.1 The Flat-Footed Response

A flat-footed attempt to square the orthodox view with *Extensionality* is superficially attractive. Recall that (A) is an ambidextrous fact, relative to PJ:

- A. Bill will give you a large sum of money to believe that he's rich.

I argued that (A) is an ambidextrous fact by drawing analogies with other non-ambidextrous facts that uncontroversially give right- or wrong-kind reasons, relative to PJ. For example, I've drawn analogies to facts like:

- E. Bill can afford to give away large sums of money without obvious personal gain.  
F. You will get a large sum of money if you believe that Bill is rich.

The flat-footed response is to say, "Hold on. (E) and (A) don't make distinct evidential contributions to the belief that Bill is rich: when I already know (E), (A) doesn't improve my evidence that Bill is rich or *vice versa*. If we just systematically weigh reasons like (E) and ignore ones like (A), then PJ avoids counterexamples from ambidextrous facts. So we can satisfy *Extensionality* by treating (A) as a wrong-kind reason --- which it is!" We can also adopt a similar approach to (B)-(D) by also treating these as wrong-kind reasons.

The flat-footed response faces at least three problems. First, it relies on the observation that the ambidextrous reason, (A) is evidentially redundant with a non-ambidextrous one, (E). But why suppose that every ambidextrous reason is similarly redundant with a non-ambidextrous one? Absent such an account, the response is insufficiently general. Indeed, recent scholarship teaches us that systematic connections between reasons are hard to find.<sup>18</sup> So we should doubt that the response generalizes adequately.

Second, the flat-footed response is *ad hoc*. (A) is clearly a right-kind or evidential reason. For example, were I to correctly base my belief that Bill is rich on (A), I would be doxastically justified or *post hoc* rational in believing that Bill is rich, absent countervailing evidence. But only beliefs based in evidence are doxastically justified or *post hoc* rational. So (A) is a right-kind reason. Likewise, when (E) isn't in my evidence, (A) raises the probability that Bill is rich, other things equal. But only evidence for *P* raises its probability. Because (A) is clearly a right-kind reason, *Extensionality* requires including (A) in PJ. Ignoring this requirement simply because doing so inconveniently undermines PJ is entirely *ad hoc*.

Finally, the response violates *Non-triviality*. *Non-triviality* forbids including or excluding a reason from a weighing analysis on the basis of facts that the analysis aims to explain. A paradigm violation of *Non-triviality* is defining the BPA's right-kind reasons as those given by the object's good-making features for goodness is what the BPA aims to analyze.

The flat-footed response does something similar. It appeals to whether a reason for belief affects the strength or degree of a belief's justification in order to determine whether it is evidence and so whether it belongs in PJ. But this appeal is viciously circular. According to PJ, evidence explains justification, not *vice versa*. Thus, excluding (A) from PJ on the grounds that it doesn't provide additional justification apart from (E) violates *Non-triviality* since the response's grounds for excluding (A) include facts about justification that are the PJ's *explanandum*.

### 5.1.2 The Sophisticated Response

Nevertheless, the flat-footed response suggests a more sophisticated one. Analyses like the BPA or PJ presuppose that some normative facts, such as facts about the competition between certain reasons, are more metaphysically fundamental than others, such as facts about value or propositional justification. Similarly, we might suppose that some normative reasons are more fundamental than others.<sup>19</sup> For example, we might imagine that (E) is a more fundamental reason to believe that Bill is rich than (A) is --- that's why their conjunction is not stronger evidence than either reason alone. The

---

<sup>18</sup> Fogal (2016) reveals fewer such connections, especially between reasons and defeaters, exist than is commonly assumed. Laskowski (2015) explores some of these connections at the level of normative thought.

<sup>19</sup> Appeal to something like fundamental reasons is not entirely uncommon. For example, Nair (2016) describes a process that he calls 'Additive Independent Reasons Accrual' that assigns a weight to a set of reasons by adding the weights of the fundamental (he calls these 'non-derivative') reasons in the set. He then goes on to express pessimism, as I also will shortly, that we can non-brutely define which reasons are fundamental. Other conceptions of fundamental reasons, such as those of Crisp (2006) and Star (2015), are incompatible with Reasons First, so they are not relevant here.

more sophisticated response to problems with ambidextrous facts asserts that only *fundamental* right-kind reasons figure in PJ and *Extensionality*. If (A) is not a fundamental reason and only fundamental right-kind reasons matter for PJ, then (A) poses no challenge to it.

An immediate problem with the sophisticated response is its dependence on the claim that (E) is more fundamental than (A) rather than *vice versa*. It is not obvious that this is true, especially to a Reasons Firster who cannot appeal to facts about propositional justification when arguing that one reason is more fundamental than the other. But more importantly than this first problem, the sophisticated response creates its own WKR problem. ‘Fundamental’, in the context of this response, plays the same role that ‘right-kind’ does in the original cases, namely, of picking out which reasons matter for the analysis. Consequently, constraints similar to those on how we understand ‘right-kind’ apply to how we understand ‘fundamental’. In particular, we need to define ‘fundamental’ in such a way that it satisfies *Non-triviality*, *Non-vacuity*, and *Extensionality* because ‘fundamental’ appears, either implicitly or explicitly, in the *analysans*.

Not only are these three desiderata a tall theoretical hurdle for the sophisticated response to clear, it’s a particularly awkward hurdle given the dialectical context. The sophisticated response we’re considering involves insisting that advocates of the orthodox view can solve the WKR problem by appealing to fundamental reasons. But the response itself creates its own WKR problem: we need to satisfy the three criteria anew, this time for ‘fundamental’ rather than ‘right-kind’, in order to advance the response. As a result, this response simply displaces the bump in the linoleum rather than smoothing it. The correct solution to a problem shouldn’t recreate the problem to be solved.

Finally, and most decisively, the response succeeds only if no ambidextrous fact gives a fundamental right-kind reason, for only then do analyses limited to fundamental reasons satisfy *Extensionality*. But some ambidextrous facts *do* give fundamental reasons. Imagine the following. You’ve promised your friend comments on her paper. But she knows that you’re forgetful, so you’ll likely forget. She devises the following mnemonic: she makes you promise to believe that you’ve made a promise in the hopes that you’ll then remember your promise to give her comments. The scenario entails that:

G. You promised to believe that you made a promise.

(G) is an ambidextrous fact relative to PJ. It gives a wrong-kind or moral reason to believe that you made a promise because there is (defeasible) moral reason to fulfil your promises and having that belief fulfils the promise. It also gives a right-kind or evidential reason to believe that you’ve made a promise because the fact represents a promise. Just as the fact that the Empire State Building is in New York is a fundamental right-kind reason to believe that the Empire State Building is in New York if anything is, the fact that you’ve made a promise is a fundamental right-kind reason to believe that you’ve made a promise, if anything is. Therefore, in any intuitive sense of the expression, (G) gives a *fundamental* right-kind reason to believe that you made a promise. So, even limiting ourselves to fundamental

right-kind reasons for belief, ambidextrous facts show that the orthodox view is incompatible with *Extensionality*, and therefore with Reasons First.

Fundamental-reason-giving, ambidextrous facts outstrip the epistemic domain.<sup>20</sup> As I suggest above, it seems plain that fundamental reasons to value an individual for their own sake flow from their character or their virtuous dispositions.<sup>21</sup> Consequently, I must locate an ambidextrous fact that represents someone's virtue since that fact gives a fundamental right-kind reason to value them. The following fact fits the bill:

H. Boris is generous.

I assume that (H) gives a fundamental right-kind reason to value Boris for his own sake. This much is relatively uncontroversial. The real work consists in showing that (H) also gives a wrong-kind reason to value him.

Irrationality merits a kind of blame. For example, when you've knowingly failed to take the necessary means to your ends, then, other things equal, you're open to a kind of criticism. As a result of this and other observations, many philosophers assume that rationality is normative. Because rationality is normative and Reasons Firsters explain normativity through claims about reasons, they find it appealing to explain instrumental irrationality as a kind of failure to respond to certain reasons.<sup>22</sup> Call these *instrumental* normative reasons.

When having an attitude is a means to one's ends, there are instrumental reasons for that attitude. For example, when your end is to get rich and you learn that having a certain belief will make you rich --- say, the belief that Bill is rich --- then facts that explain how having that attitude will help make you rich give instrumental reasons for the belief. Likewise, when your end is to submit certain immigration documents and you learn that only by fearing Ba'al-zebul will you be able to submit them, then facts that explain why the fear is necessary for submitting the documents will give instrumental reasons for it. Moreover, since these instrumental reasons do not help to make those attitudes fitting, they are wrong-kind reasons for those attitudes.

As a result, (H) is a wrong-kind reason for someone to value Boris when it helps to explain why valuing Boris furthers that person's ends. In such a circumstance, just as there are wrong-kind reasons to fear

---

<sup>20</sup> Thanks to a referee for *Philosophers' Imprint* for pushing me towards a clearer rebuttal of the sophisticated response.

<sup>21</sup> Relying on character in what follows is perhaps vulnerable to worries from the so-called 'situationist' challenge to the idea that human beings have dispositions that approximate anything like character traits -- see, for example, Harman (1999). However, I'm not committed to the existence of such traits. I'm merely positing them on behalf of my opponent because *if* some reasons to value people are fundamental, then those reasons follow from virtuous character traits or dispositions. If there are no such traits, then the fundamental/derivative distinction does not plausibly apply to normative reasons to value people. If that's so, then the objection I rebut above does not achieve liftoff.

<sup>22</sup> See Schroeder (2007a) and Lord (2018) for some developed versions of this position.

Ba'al-zebul or to believe that Bill is rich when it furthers one's ends, (H) gives a wrong-kind reason to value Boris. Here's one such scenario: suppose that Selfish Sam is an ambitious intern looking for a raise at a company where only two people can grant it: Boris and Mr. Miser. While researching the best strategy to get that raise, Sam learns that generous people benefit those who value them for their own sake. Since Sam knows that Boris is generous and that Mr. Miser is not, (H) gives Sam a reason Sam to value Boris for his own sake in order to get the raise.<sup>23</sup> That makes (H) an instrumental reason to value Boris, which is a wrong-kind reason to do so. So (H) is an ambidextrous fact that gives a fundamental right-kind reason and a wrong-kind reason for Sam to value Boris.

It's especially controversial which moral reasons are fundamental (if any are), so I won't rebut a parallel defense of the orthodox view and MR. Any example I present will be highly tendentious. Nevertheless, (G) and (H) make it extremely likely that ambidextrous facts undermine even a version of MR restricted to fundamental reasons, especially since I've just shown how to construct the rebuttal using instrumental reasons given any conception of fundamental moral reasons. As a result, appeals to fundamental right-kind reasons generally fail to exclude ambidextrous facts from weighing analyses.

## 5.2 Defending Premise (2)

The other way to resist the argument is to deny that Reasons First implies *Extensionality*. I've heard several versions of this response from various philosophers. However, what they usually fail to recognize is that it dispenses with *Extensionality* only by drawing the right-/wrong-kind distinction among entities that are not reasons. Doing so may solve problems with ambidextrous facts but it requires using objects that are not reasons in analysis. Consequently, the approach is not compatible with Reasons First. I discuss two such proposals to give you their general flavour.

### 5.2.1 Weighing Things Other Than Reasons

According to one response, we must distinguish the different weights associated with an ambidextrous fact.<sup>24</sup> For example, (A) has a right-kind or evidential weight and a wrong-kind or non-evidential weight. According to the proposal, these weights are our *analysans*: weights are of the right- or wrong-kind for the analysis of a normative property and their competition, not the competition between facts, underpins normative analysis. To analyze propositional justification, for example, we consider whether (A)'s right-kind or evidential weight is sufficient, ignoring its wrong-kind weight. Distinguishing weights quarantines (A)'s malign pragmatic influence on PJ. As a result, we can

---

<sup>23</sup> It might be observed that if Sam values Boris in order to get a raise, Sam does not value Boris for his own sake. That shows only that Sam is going to have a hard time getting the raise if he continues to be selfish. It does not show that (H) is not a wrong-kind reason to value Boris for his own sake. Moreover, (H) clearly *can* motivate someone to value Boris for his own sake --- after all, on orthodox assumptions, it's a right-kind reason to value Boris for his own sake. Consequently, it survives challenges to the existence of wrong-kind reasons that insist that normative reasons must be motivating reasons, such as from Hieronymi (2005) and Rowland (2015).

<sup>24</sup> Joe Horton, Jonathan Way, and Tristram McPherson each suggested that I distinguish weights to solve the problem on separate occasions.



dispense with *Extensionality*; we needn't invoke a distinction between right- and wrong-kind reasons now that one obtains between weights:

*Propositional Justification\** (PJ\*): A belief is *propositionally justified* just in case and because there is evidentially sufficient weight to have it.

The idea of weighing weights is odd. We time moments using seconds and measure lengths using metres. We do not time seconds or measure metres. These banalities suggest that, just as we weigh luggage not pounds, we weigh reasons not their weights. But despite its oddity, the proposal allows us to analyze normative properties like propositional justification without relying on *Exclusivity*.

However, even if PJ\* is true, it does not reconcile the orthodox view of reasons with Reasons First. Recall reasons' characteristic role, offered in section two:

*Reason*: An object *o* is a normative reason iff *o* is the kind of thing weighed in a weighing explanation of a normative property.

PJ\* does not use facts in its weighing explanation of propositional justification. Therefore, given *Reason*, PJ\* implies that facts are not reasons. Likewise, it does use weights -- presumably properties of facts -- as the thing weighed in weighing explanations. Weights are not facts. As a result, the solution implies that something *other* than facts are reasons, so the solution does not reconcile Reasons First with the orthodox view of reasons.<sup>25</sup>

### 5.2.2 Adding to the Reason Relation

When something is a reason for action, it favours *someone* doing *something*.<sup>26</sup> The many-place favouring relation expressed in the preceding sentence is the 'reason relation', which is popularly thought to hold between (at least) agents and options like actions, beliefs, feelings, or outcomes.<sup>27</sup> I'll represent the favouring relation using an ordered triple,  $\langle(\text{reason}), (\text{agent}), (\text{option})\rangle$ , though the relation may in fact contain more arguments.

---

<sup>25</sup> Other apparent solutions structurally resemble the weight proposal. For example, we might distinguish say (A) gives (A)-*qua*-right-reason and (A)-*qua*-wrong-reason and weigh some *qua*-reasons rather than others. This proposal accommodates ambidextrous facts, but it is incompatible with the orthodox view. Since (A)-*qua*-right-reason and (A)-*qua*-wrong-reason are distinct, they are not identical, and, therefore, (A) is not identical with both. So *qua*-reasons are incompatible with the orthodox view.

<sup>26</sup> Nagel (1970) argues that some reasons, agent-neutral reasons, lack an argument space for agents. It is now popular to model agent-neutral reasons as universally quantifying over the domain of agents. I assume the latter approach.

<sup>27</sup> Precisely which arguments form the reason relation is a matter of controversy. For example, some hold that *outcomes* or *answers to questions* not *actions* are the options favoured by reasons. Others posit more arguments. My arguments are independent from this debate. For more, see Howard and Schroeder (fc, Ch.3&4).

The second proposal recommends that we add a new parameter to the reason relation that tracks the difference between, for example, moral and non-moral reasons for action.<sup>28</sup> I'll use 'M' and 'N', respectively, as possible values for this parameter without offering a concrete proposal for what they represent. On this conception of the reason relation, when (D) --- that Richie Rich is drowning --- gives both a moral and a non-moral reason to save Richie, that's in virtue of standing in two different reason relations. It gives a non-moral reason when it stands in  $\langle (D), \text{you, to save Richie, N} \rangle$  and a moral one when it stands in  $\langle (D), \text{you, to save Richie, M} \rangle$ .

There are some immediate worries with the proposal. First, it is *ad hoc*. Second, it is hard to imagine a value for P and M that is consistent with Reasons First. But third, and most troublingly, even if we add a new parameter to the reason relation, we exclude (D)'s non-moral influence using this new parameter only if we also add it to MR:

*Moral Requirement\** (MR\*): An act is *morally required* just in case and because there is a conclusive moral reason to do it, *relative to M*.

MR\* is consistent with Reasons First only if it analyzes moral requirements strictly in terms of facts about the competition between moral reasons. But MR\* appeals to more than reasons; it appeals to facts about that competition *relative to a further parameter*, M. M, whatever it is, therefore does serious analytical work in MR\* by distinguishing a fact's right- or wrong-kind bearing. Because MR\* appeals to more than facts about the competition between reasons, it is incompatible with Reasons First. So MR\* does not reconcile Reasons First with the orthodox view of reasons.

## 6. A New Hope

The argument that I've offered is simple. Reasons First implies *Exclusivity*, which implies that whatever reasons are, they cannot straddle right- and wrong-kind categories. Facts like (A)-(G) straddle these categories. Therefore, if Reasons First is true, reasons cannot be numerically identical to those facts. Reasons Firsters must reach for a different account of the substance of reasons.

Subtle variations on the orthodox view exist. For example, some think that each normative reason is a proposition, even sometimes a false one (*e.g.*, the Mark Schroeder of *Slaves of the Passions*, Juan Comesaña, and Jeremy Fantl). Others think that reasons are states-of-affairs (*e.g.*, Jonathan Dancy). However, the argument from ambidextrous facts generalizes to these views as well, showing that they too are incompatible with Reasons First. The argument generalizes to these views because all three share an assumption about what it is to give a reason: that when a consideration gives a reason, it is numerically identical to that reason. The problem posed by ambidextrous facts is rooted in this assumption. For example, (A) clearly gives a right-kind reason, RKR, and a wrong-kind reason, WKR. Given the identity assumption, (A)=RKR and (A)=WKR. The transitivity and symmetry of identity

---

<sup>28</sup> For example, the view defended in Finlay (2014) plausibly includes *ends* in the reason relation.

implies the fatal result that  $RKR=WKR$ . Thus, regardless of whether reasons are facts or propositions or states-of-affairs, Reasons First cannot satisfy *Extensionality* if this implication is true.

Addressing the problem that ambidextrous facts create for Reasons First therefore requires rethinking what it means for a fact to ‘give’ a reason. Philosophers have problematically assumed that the relation is identity. But there are alternatives. For example, we might suppose that reasons are not identical with facts but constituted by them in the way that a statue is constituted by its clay or a conjunction is constituted by its conjuncts. Constitution, unlike identity, is anti-symmetric. As a result, even if (A) constitutes both a RKR and a WKR for some attitude, we cannot derive that  $WKR = RKR$  from that fact. The thought that certain facts *constitute* reasons therefore does not entail violations of *Exclusivity*.

This alternative conception of the ‘giving’ relation between facts and reasons raises its own questions. In particular, it raises the question of how RKR and WKR differ when they are given by the same fact, which is the kind of question that dominates discussion concerning the WKR problem. Nevertheless, I conjecture that thinking of the relation between facts and reasons along the lines of constitution not identity, as the orthodox account implies, promises new resources for solving the WKR problem.<sup>29</sup>

---

<sup>29</sup> Indeed, in Howard (fc-a), Howard (fc-b), and Howard (fc-c), I make steps towards advancing an alternative account of the substance of reasons according to which reasons are only partly constituted by facts.

## REFERENCES

- Broome, John. 2004. "Reasons." In R. Jay Wallace (ed.), *Reason and Value: Themes From the Moral Philosophy of Joseph Raz*. Oxford University Press. pp. 2004--28.
- Broome. 2013. *Rationality Through Reasoning*. Oxford: Wiley.
- Crisp, Roger. 2000. "Review of 'Value ... And What Follows' By Joel Kupperman." *Philosophy* 03: 452--62.
- Crisp, Roger. 2005. "Value, Reasons and the Structure of Justification: How to Avoid Passing the Buck." *Analysis* 65 (1):80--85.
- Crisp, Roger. 2006. *Reasons and the Good*. Oxford: Oxford University Press.
- Dancy, Jonathan. 2000. *Practical Reality*. Oxford: Oxford University Press.
- Danielsson, Sven and Jonas Olson. 2007. "Brentano and the Buck-Passers." *Mind* 116 (463):511 - 522.
- D'Arms, Justin and Daniel Jacobson. 2000. "Sentiment and Value." *Ethics* 110: 722-748.
- Finlay, Stephen. 2014. *A Confusion of Tongues*. Oxford University Press.
- Fogal, Daniel. 2016. "Reasons, Reason, and Context." In Errol Lord & Barry Maguire (eds.), *Weighing Reasons*. Oxford: Oxford University Press.
- Gert, Joshua. 2016. "A Fitting End to the Wrong Kind of Reason Problem." *Ethics* 126 (4):1015-1042.
- Harman, Gilbert. 1999. "Moral Philosophy Meets Social Psychology." *Proceedings of the Aristotelian Society* 99 :315-331.
- Heuer, Ulrike. 2010. "Beyond Wrong Reasons: The Buck-Passing Account of Value." In *New Waves in Metaethics*. Edited by Michael Brady. Palgrave-Macmillan.
- Hieronymi, Pamela. 2005. "The Wrong Kind of Reason." *The Journal of Philosophy* 102 (9): 437--57.
- Hieronymi, Pamela. 2013. "The Use of Reasons in Thought (and the Use of Earmarks in Arguments)." *Ethics* 124 (1):114-127.
- Howard, Christopher. 2016. "In Defense of the Wrong Kind of Reason." *Thought: A Journal of Philosophy* 5 (1): 53-62. 2016.

- Howard, Nathan Robert. Forthcoming-a. "The Goals of Moral Worth." *Oxford Studies in Metaethics Vol. 16*. Edited by Russ Shafer-Landau. Oxford: Oxford University Press.
- Howard, Nathan Robert. Forthcoming-b. "Primary Reasons as Normative Reasons." *Journal of Philosophy*.
- Howard, Nathan Robert. Forthcoming-c. "One Desire Too Many". *Philosophy and Phenomenological Research*.
- Howard, Nathan Robert, and Schroeder, Mark. Accepted. *The Fundamentals of Reasons*. Oxford University Press.
- Kearns, Stephen, and Daniel Star. 2015. "Weighing Explanations." In *Weighing and Reasoning: Themes from the Philosophy of John Broome*. Edited by Iwao Hirose and Andrew Reisner. Oxford: Oxford University Press.
- Laskowski, Nicholas. 2015. "Non-Analytical Naturalism and the Nature of Normative Thought: A Reply to Parfit." *Journal of Ethics and Social Philosophy* (1):1-5.
- Leary, Stephanie. 2017. "In Defense of Practical Reasons for Belief." *Australasian Journal of Philosophy* 95 (3):529-542.
- Liao, S. Matthew. 2010. "The Buck-Passing Account of Value: Lessons From Crisp." *Philosophical Studies* 151 (3):421 - 432.
- Lord, Errol. 2018. *The Importance of Being Rational*. Oxford University Press.
- Maguire, Barry, and Snedegar, Justin. forthcoming. "Normative Metaphysics for Accountants" *Philosophical Studies*.
- Nagel, Thomas. 1970. *The Possibility of Altruism*. Princeton University Press.
- Nair, Shyam. 2016. "How Do Reasons Accrue?" In *Weighing Reasons*. Edited by Errol Lord and Barry Maguire. Oxford: Oxford University Press.
- Olson, Jonas. 2004. "Buck-Passing and the Wrong Kind of Reasons." *Philosophical Quarterly* 54: 295-300.

- Parfit, Derek. 2001. "Rationality and Reasons." In *Exploring Practical Philosophy*. Edited by Dan Egonsson et al. Burlington, VT: Ashgate.
- Piller, Christian. 2006. "Content-Related and Attitude-Related Reasons for Preferences." *Royal Institute of Philosophy Supplements* 81 (59): 155–81.
- Rabinowicz, Wlodek, and Toni Rønnow-Rasmussen. 2004. "The Strike of the Demon: On Fitting Pro-Attitudes and Value." *Ethics* 114 (3): 391–423.
- Raz, Joseph. 2011. *From Normativity to Responsibility*. Oxford: Oxford University Press.
- Reisner, Andrew. 2009. "The Possibility of Pragmatic Reasons for Belief and the Wrong Kind of Reasons Problem." *Philosophical Studies* 145 (2):257 - 272.
- Ross, W.D. 2002. *The Right and the Good*. Edited by Philip Stratton-Lake. Oxford: Clarendon Press.
- Rowland, Richard. 2015. "Dissolving the Wrong Kind of Reason Problem." *Philosophical Studies* 172 (6):1455-1474.
- Scanlon, T. M. 2001. *What We Owe to Each Other*. Cambridge, MA: Harvard.
- Schroeder, Mark. 2007a. *Slaves of the Passions*. Oxford: Oxford University Press.
- Schroeder, Mark. 2007b. "Weighting for a Plausible Humean Theory of Reasons." *Nous* 41 (1): 138-160.
- Schroeder, Mark. 2010. "Value and the Right Kind of Reason." *Oxford Studies in Metaethics* 5:25-55.
- Schroeder, Mark. 2018. "Getting Perspective on Objective Reasons" *Ethics* 128 (2).
- Schroeder, Mark. Forthcoming. *Reasons First*. Oxford University Press.
- Sharadin, Nathaniel. 2013. "Schroeder on the Wrong Kind of Reasons Problem for Attitudes" *Journal of Ethics and Social Philosophy* 7 (3): 1-8. 2013.
- Star, Daniel. 2015. *Knowing Better*. Oxford: Oxford University Press.

Stratton-Lake, Philip. 2011. "Eliminativism about Derivative Prima Facie Duties." In *Underivative Duty: British Moral Philosophers from Sidgwick to Ewing*. Edited by Thomas Hurka. Oxford: Oxford University Press.

Way, Jonathan. 2012. "Transmission and the Wrong Kind of Reason." *Ethics* 112 (3): 489-515.