

## In Defense of the Right Kind of Reason

Chris Howard and Stephanie Leary

(forthcoming in *Fittingness* volume by Oxford University Press, edited by Chris Howard and Richard Rowland)

Why do you love your child? Understood as a request for justification, it's natural to give two different kinds of answers. On the one hand, you might list some of her lovable features: she's funny, intellectually curious, kind, and spirited. On the other hand, you might list some ways in which loving her does some good, e.g., by promoting her happiness and by keeping you from ditching her on the side of the road when she's driving you crazy. The former answer appeals to *right-kind* reasons (RKR): considerations to do with whether your child is lovable and thus fitting for you to love. The latter answer appeals to *wrong-kind* reasons (WKR): considerations to do with the goodness of loving your child, whether or not your love for her is fitting. This distinction generalizes for other attitudes like admiration, desire, and belief. Considerations to do with whether something is admirable, desirable, or credible constitute RKRs for admiring, desiring, and believing (respectively) because they contribute to the fittingness of these attitudes. Whereas considerations to do with whether it's good to admire, desire, or believe something constitute WKRs for admiring, desiring, or believing.<sup>1</sup>

Since the origin of this distinction, there have been WKR skeptics: folks who claim that only RKRs genuinely favor, or justify, having an attitude. For example, these skeptics insist that the only facts relevant to whether you should love someone are those relevant to whether love is fitting. The fact that loving someone would do some good at best bears on whether you should want to love them (as an RKR for desiring) or on whether you should try to bring it about that you do.<sup>2</sup>

Recently, though, many WKR-enthusiasts have become RKR skeptics: they claim that WKRs are the only facts that genuinely favor, or justify, having an attitude and that putative RKRs for attitudes are merely formally normative reasons that arise from attitudes' constitutive standards of correctness.<sup>3</sup> Constitutive standards of correctness lack normative authority. The reasons arising from such standards don't bear on what you really ought to or may do, unless you've got authoritative reason to be engaged in whatever activity the standards govern. RKR skeptics thus insist that the fact that someone has lovable properties can at best be derivatively relevant to whether you really ought to love them: a person's lovable properties can be relevant to whether you should love them only if you've got authoritative reason to love *correctly*, i.e., according to the standards internal to love. Your daughter's lovable properties thereby lack a kind of normative authority: they fail to provide non-derivatively authoritative normative reasons for love, in the sense that they can't, by themselves, contribute to making it the case that you really ought to, or may, love your daughter.<sup>4</sup>

---

<sup>1</sup> For a helpful survey of the literature on the RKR/WKR distinction, see (Gertken and Kiesewetter 2017).

<sup>2</sup> See, e.g., (Parfit 2011), (Skorupski 2010), and (Way 2012).

<sup>3</sup> See esp. (Côté-Bouchard and Littlejohn 2018), (Mantel 2019), and (Maguire and Woods 2020).

<sup>4</sup> On the distinction between "authoritative" and "merely formal" normativity, see esp. (McPherson 2018), (Wodak 2019), and (Woods 2018). We remain neutral about how exactly this distinction should be drawn, but we take the idea to be intuitive: something's normativity is authoritative when it's relevant to settling what you "genuinely" or "really" ought to or may do. Its authoritative normativity is "non-derivative" if it doesn't depend on the authoritative normativity of some other factor.

We reject both kinds of skepticism: we think both RKR and WKR are authoritatively normative. Elsewhere, we've defended the authoritative normativity of WKRs.<sup>5</sup> Here, we defend the authoritative normativity of RKRs. Our main claim is that RKRs are authoritatively normative in the sense that they play a non-derivative role in weighing explanations of an attitude's overall, authoritative deontic status. This leaves open the question of whether or how RKRs compare with, or weigh against, WKRs.<sup>6</sup> It's also neutral about whether RKRs ever demand, rather than merely justify, the attitudes they favor, i.e., whether RKRs are requiring or only permissive.<sup>7</sup>

In §1, we argue that the intuitive data about cases favors the view that RKRs are authoritatively normative. We then argue that the kinds of theoretical considerations that RKR skeptics appeal to in order to motivate their skepticism can and should be resisted. In §2, we address the most common theoretical motivation for RKR skepticism, viz., the assumption that RKRs arise from standards of correctness internal to the attitudes. We argue that fittingness is distinct from constitutive correctness and that once we appreciate that RKRs are fit-makers, rather than constitutive-correctness-makers, RKR skepticism is no longer warranted. Then, in §3, we respond to Barry Maguire's (2018) recent argument that fit-making facts can't be normative reasons at all since they lack certain essential features of reasons. We argue that fit-making facts have the features of reasons that Maguire claims they lack.

## 1. Intuitive data

The main reason to believe that RKRs are authoritatively normative is that this claim best explains intuitions about cases. Consider cases where having an attitude would have no value at all, or even a little disvalue, but there are facts that make the attitude fitting. For example, suppose you meet someone at a party who has many admirable features (they're intelligent, accomplished, they dedicate themselves to philanthropic pursuits, and they're sincerely warm and kind), but your privately admiring them would have no good effects on you, them, or anyone else. Still, it seems you should admire this person.<sup>8</sup> Or suppose that while you're waiting in line to see a movie, an audience member from an earlier showing walks by and blurts out the surprise ending. Believing the fellow movie-goer's testimony would be bad for you since it makes the movie less enjoyable. But it still seems like you should believe it.<sup>9</sup> The authoritative normativity of RKRs explains these verdicts: if fit-making facts bear non-derivatively on whether you should have some attitude, this explains why you should admire someone who's admirable or believe something that's credible even if doing so would have no benefit or would be slightly harmful. RKR skeptics, on the other hand, must reject these verdicts. We think this is a real cost. These sorts of intuitions are the basic data that any normative theory should explain.<sup>10</sup>

---

<sup>5</sup> See (Howard 2016, 2019) and (Leary 2017).

<sup>6</sup> On which see (Berker 2018), (Howard 2019b), and (Reisner 2008, ms).

<sup>7</sup> For discussion, see (Whiting 2021) and (Berker this volume).

<sup>8</sup> We use "should" here in the authoritatively normative sense, but in a way that's neutral with respect to whether RKRs require, permit, or justify attitudes in a sense that's stronger than permission, but weaker than requirement (see [Berker this volume] for discussion of this latter possibility).

<sup>9</sup> This case comes from (Kelly 2003).

<sup>10</sup> One way in which RKR skeptics might try to capture this data is by claiming that whenever there's an RKR for an attitude, there's a WKR for the attitude with equal weight because it's always in some way good to have an attitude for which there are RKRs. We lack the space to discuss this strategy here, but for reasons to

Can we generalize from cases like the above to the claim that *all* RKR are authoritatively normative? We think so, unless there are specific examples to the contrary. Richard Rowland (fc) suggests several examples of attitudes the RKR for which might seem to lack normative authority: (a) highly specific attitudes like *schadenfreude* or *chrysalism* (the tranquil feeling of being indoors during a storm), (b) vicious attitudes like envy or anger, and (c) boredom and depression. But we argue here that none of these examples suggest that some RKR aren't authoritatively normative.<sup>11</sup>

Consider first highly specific attitudes like *schadenfreude* and *chrysalism*. Rowland seems to take such attitudes to show that there's a vast and incredibly diverse array of attitudes, each of which can be fitting with respect to a highly specific object. And Rowland thinks it would be odd to think that there are thereby authoritative reasons to have such attitudes toward those objects.

But we think that it's a mistake to think that attitudes like *schadenfreude* and *chrysalism* are distinctive attitudes that have their own fittingness conditions. Rather, they're ordinary, run-of-the-mill attitudes toward very specific objects: *schadenfreude* is *being pleased* at the suffering of others; *chrysalism* is *being relieved* at being indoors during a storm. So, when we ask whether *schadenfreude* or *chrysalism* can be fitting, we're asking whether it can be fitting to be pleased by the suffering of others, or to be relieved about being indoors in a storm. These are substantive normative questions that aren't settled by the mere fact that *schadenfreude* and *chrysalism*, as a matter of conceptual definition, involve having these attitudes toward those objects. The claim that all RKR are authoritatively normative thus doesn't imply that there are authoritative reasons for attitudes like *schadenfreude* and *chrysalism*. Whether there are depends on whether these attitudes can be fitting.

Next consider vicious attitudes—envy, in particular.<sup>12</sup> Rowland (fc: 9) claims that although envy can be fitting when directed at the enviable, if one fails to envy the enviable one isn't normatively criticizable. He takes this to show that RKR for envy aren't authoritatively normative.

We have three responses. First, from the fact that envy is fitting toward the enviable it doesn't follow that there are RKR for envy. To see this, consider that for any attitude that can be fitting, it's an open question what properties comprise the evaluative property to which the fittingness of the attitude corresponds. Similarly, it's an open question whether that property can ever be instantiated. For example, in the case of envy, it's an open question what properties comprise *enviability* and whether this property can be instantiated. Suppose that for someone to be enviable is for them to possess something desirable that you lack, where this difference in possession itself is bad for you.<sup>13</sup> It's highly controversial whether the mere fact that someone has something desirable that you lack could ever be bad for you. Hence, it's unclear whether anyone can ever be enviable. If not, then it can't be fitting to

---

reject it; see, e.g., (Howard 2019a, 2019b) and (Way 2013). It's worth noting that many RKR skeptics also reject this strategy; see, e.g., (Maguire and Woods 2020), (Papineau 2013), and (Rinard 2019).

<sup>11</sup> Lord and Sylvan (2019) also argue that not all RKR are authoritatively normative by appealing to examples of RKR for action. We don't address their examples here because we think their examples rest on the mistaken assumption that RKR arise from constitutive standards of correctness, which we address in §2.

<sup>12</sup> We focus on envy because anger is more contested; see esp. (Cherry 2018) and (Srinivasan 2017).

<sup>13</sup> This account is suggested by D'Arms and Jacobson (2006).

envy anyone, and so, there can't be RKR for envy. But this is very different from the claim that there are RKR for envy, but that they lack normative authority.<sup>14</sup>

Second, even assuming that someone could be enviable and thus that there could be RKR for envy, there are possible explanations, compatible with the authoritative normativity of RKR, for why one needn't be normatively criticizable for failing to envy the enviable. One such explanation is that some (or all) RKR might not be requiring: if the fact that someone is enviable bears on whether you should envy them in a sense that is weaker than requirement, then you wouldn't be criticizable if you fail to envy them despite the RKR to do so.<sup>15</sup> Another possible explanation is that there may be strong WKR against envy in general if envying is always bad for the envier or morally vicious (or both). So, even if someone is truly enviable such that there are authoritatively normative RKR to envy them, it may nonetheless be the case that you ought not to envy them, since there are stronger WKR against your doing so.

Third, and finally, the claim that there are authoritatively normative RKR for envy seems plausible if we compare envying the enviable against envying the unenviable. Other things being equal, it seems like there's less (authoritatively normative)<sup>16</sup> reason for someone to envy the unenviable than there is for someone to envy the enviable, which suggests that there's some reason to envy the enviable.<sup>17</sup> Of course, the claim that there's less reason to do A than there is to do B doesn't immediately entail that there's some reason to do B; there could be less total reason to do A than to do B only because there's more reason not to do A than there is not to do B. But that doesn't seem like a viable explanation in the present case: why think that there's less reason against envying the enviable than there is against envying the unenviable? So instead, what explains the fact that there's less reason to envy the unenviable than there is to envy the enviable must be that there's some positive, authoritative RKR for the latter.

Rowland's final examples concern boredom and depression. Again, Rowland (fc: 3-4) claims that we aren't criticizable or "normatively missing something" if we're never bored or depressed and that this suggests that RKR for boredom and depression aren't authoritatively normative.

We think it's important to distinguish between object-directed attitudes, like belief and admiration, and mental states that aren't object directed, like mere feelings or moods. We think there are RKR only for the former. So when Rowland discusses boredom and depression, we think it's important what kinds of mental states he takes these to be. If he's talking about boredom and depression as moods, for example, then we think there are no RKR for these mental states, and that this is what explains why one needn't be normatively criticizable if one is never in them.

---

<sup>14</sup> Sara Protasi (this volume) defends the claim that envy can be fitting (and hence that the property of being enviable can sometimes be instantiated). By our lights, this defense doubles as a defense of the claim that there are RKR for envy that and these reasons are authoritative.

<sup>15</sup> This explanation is compatible with RKR being reasons that favor or justify the attitude in a sense that's stronger than permission, but weaker than requirement (as Berker [this volume] suggests).

<sup>16</sup> We omit this clarification from now on, but in what follows it should be understood that we use 'reason' to refer to authoritatively normative reason.

<sup>17</sup> For arguments of this general form in favor of establishing the existence of a positive pro tanto reason for an attitude, see, for example, (Schroeder 2007) and (Wodak fc).

But some instances of boredom and depression do seem like object-directed attitudes—for example, if you’re bored by a talk or depressed about a pandemic. If Rowland is talking about these kinds of boredom and depression, then what we said above about envy applies here too. First, it’s an open question whether these attitudes can ever be fitting and thus whether there can be RKR for them at all. Perhaps they can—perhaps it’s fitting to be bored by the boring and depressed by the depressing and these are both evaluative properties that can be instantiated. But, second, there are alternative explanations for why one needn’t be criticizable if one fails to be bored by a boring talk or depressed by a depressing pandemic: RKR may not be requiring, or there may be strong WKR not to be bored or depressed. And, finally, if we compare being bored by the boring or being depressed by the depressing to being bored by the exciting and being depressed by the joyous, there seems to be less authoritative reason to be bored and depressed in the latter cases than there is in the former.

In sum, then, we think intuitive verdicts about certain cases can be explained only if RKR are authoritatively normative and that the intuitive verdicts about Rowland’s cases can be explained in a way that’s consistent with taking all RKR to be authoritatively normative. Some RKR skeptics like Rinard (2019) and Papineau (2013) argue that we should accept the view that only WKR are authoritatively normative because it’s simpler and more unified. But we think these virtues provide good enough reason to prefer one theory to another only if the former is equally capable of explaining all the intuitive data. And RKR skeptics can’t explain it. This is a serious cost, which we should accept only if there are further, strong theoretical reasons for thinking that, despite the appearances, RKR aren’t authoritatively normative after all. So, we turn now to some arguments that purport to establish this.

## **2. Right-kind reasons as constitutive-correctness-makers**

The main theoretical argument against the authoritative normativity of RKR starts with the claim that RKR arise from constitutive standards of correctness. On this view, for a fact R to be an RKR for an attitude A is for R to contribute toward making it the case that A is correct according to a standard that’s constitutive of the kind of attitude that A is. But constitutive standards of correctness aren’t authoritatively normative. From the fact that a chess move would be correct according to the standards constitutive of chess playing, nothing immediately follows about whether we genuinely ought to, may, or even have reason to, make it. We have reason to make constitutively correct chess moves only if we have reason to be playing chess (and trying to win) in the first place. Indeed, as this example makes clear, constitutive standards of correctness come cheap: we can create them out of thin air by inventing new games, activities, clubs, etc. But authoritative normativity isn’t cheap in this way. So, if RKR are constitutive-correctness-makers, then they also lack normative authority.<sup>18</sup>

One response would be to argue that at least some constitutive standards of correctness are authoritatively normative.<sup>19</sup> Metaethical constitutivists might be sympathetic to this, but we’re not. We tend to agree with our objector that constitutive correctness isn’t authoritatively normative. Our preferred response is to reject that RKR should be analyzed in terms of constitutive correctness. We argue that not only is this analysis not mandatory, but there are good reasons to resist it.

---

<sup>18</sup> See esp. the references in note 3.

<sup>19</sup> This is Lord and Sylvan’s (2019) response.

Start with some history. The terminology of “right-kind reasons” has its source in the literature on the “wrong kind of reason problem” for “buck-passing” analyses of value. Buck-passing analyses analyze various evaluative properties in terms of reasons for various sorts of response. For example, they might claim that what it is for something to be *admirable* is for there to be (sufficient) reason to admire it. The WKR problem is the problem that such analyses seem subject to counterexamples: I might have reasons to admire something that isn’t admirable; perhaps I’ll get the goods if I do. If these reasons really are reasons to admire, then some reasons—like these—are of the “wrong kind” to figure in buck-passing accounts. Right-kind reasons are of the “right kind” to figure in buck-passing accounts. What are the right kinds of reasons to figure in a buck-passing account of admirability? The answer: all and only those reasons to admire that can contribute to making it the case that the object of admiration is *admirable*. RKR for admiration are thus facts that ensure the existence of the relation to admiration that it takes for its object to be admirable. We have a name for this relation: it’s ‘fittingness’.

The fittingness relation is the relation in which a response stands to its object when the object merits—or is worthy of—that response. RKR for an attitude are thus those that are relevant to whether its object is worthy of it. This checks out. For something to be admirable is for it to be worthy of admiration, for something to be desirable is for it to be worthy of desire, and so on. Thus, RKR for an attitude *A* are those that can contribute to making it the case that *A*’s object merits, or is worthy of, the kind of attitude *A* is. In short: RKR for an attitude are those that can contribute to explanations of the attitude’s fittingness—they’re fit-making facts for the attitude in question.<sup>20</sup>

The upshot is this: RKR can be plausibly analyzed in terms of constitutive correctness only if fittingness can. Since RKR are facts that make responses fitting, they can be correctness-makers only if fit-makers are. So, is it true that what it is for an attitude to be fitting, or merited, is for it to be correct according to a standard that’s constitutive of the kind of attitude in question? In the last decade this suggestion has become quite popular, but we’re unconvinced.

The first thing to note is that, despite its current popularity, the proposed analysis is revisionary. Although fittingness has only recently come to occupy the limelight in 21<sup>st</sup> century normative philosophy, the relation has a long history. For example, it features prominently in the work of many normative theorists writing in the late nineteenth and mid-twentieth centuries (Brentano 1889/2009, Brandt 1946, Broad 1930, Ewing 1947,) as well as in the work of slightly more recent writers (Feinberg 1970, Gibbard 1990, McDowell 1998, Wiggins 1987).<sup>21</sup> But none of these authors holds that the fittingness of an attitude is a matter of its satisfying a constitutive standard of correctness.<sup>22</sup> As far we can tell, this idea has its origins in a relatively recent paper by Mark Schroeder

---

<sup>20</sup> The terminology of ‘fit-making’ isn’t perfect since ‘making’ is a success verb and, as we’ll explain below (in sect. 3), there can be fit-making facts for attitudes that aren’t fitting. A better term would be ‘fit-contributors’, where fit-contributors stand in a non-factive ‘contributing-to-making-it-the-case relation’. This clarification is prompted by Kiesewetter (2021), who makes a parallel point regarding reasons as ‘justification-makers’.

<sup>21</sup> See (Howard 2018: 2n4).

<sup>22</sup> Brentano uses ‘correctness’ [*Richtigkeit*] to refer to the relation we’re calling ‘fittingness’, but doesn’t understand an attitude’s correctness in terms of its satisfying a norm that’s constitutive of it. To anticipate: we’re happy to call an attitude ‘correct’ if it’s fitting, as long as the relevant kind of correctness isn’t constitutive correctness, but rather correctness according to a norm that’s external to the attitude. More on this below.

(2010). Since Schroeder's paper, many others have followed suit.<sup>23</sup> Indeed, in a forthcoming paper in which he argues *for* the authoritative normativity of RKR's, Benjamin Kiesewetter remarks that "it seems fairly uncontroversial to say that the right kind of reasons for attitudes are essentially linked to the constitutive correctness standards for the attitude in question" (2021: 2).<sup>24</sup>

But there are reasons to doubt this recently popular view. First, in the case of action, fittingness and constitutive correctness come apart. If you strap on your skates and perform a triple axel in the middle of your department meeting, then you might perform the jump correctly, but the jump (at least in normal circumstances) wouldn't be merited. A response is merited only when it's in some sense "called for" by the situation, or by certain specific features of it. But there's no sense in which your correct action—the triple axel—is called for in the situation under consideration.

Consider also assertion. Some hold that assertion has a constitutive standard such that, by nature, assertions are correct if and only if they're true.<sup>25</sup> On this view, what makes an assertion correct, when it is, is the fact that it's true. But an assertion isn't merited in virtue of being true. Many trivial, private, or impolite claims, though true, don't merit being asserted in most contexts. And some claims that merit assertion—e.g., "I appreciate you helping me move"—may be merited not because they're true but in virtue of something else, e.g., being an expression of gratitude.<sup>26</sup> This suggests that, insofar as fittingness is meritedness or worthiness, fittingness and correctness aren't the same.

The proponent of fittingness-as-constitutive-correctness might respond by conceding that fitting action isn't necessarily correct action and simply restrict their thesis to the view that what it is for attitudes, specifically, to be fitting is for them to be correct according to their constitutive standards. We find this reply unattractive because it proposes that fittingness isn't a unified category, that the fittingness of an action is a very different property than the fittingness of an attitude. In our view, this is a mistake. Considerations of theoretical unity, simplicity, etc. suggest that the default view should be that 'fitting', as it might be properly predicated of actions, refers to the same property that it does when properly predicated of attitudes. The proponent of fittingness-as-constitutive-correctness who proposes to restrict their thesis to attitudes therefore owes us some rationale for thinking that the fittingness of attitudes can be analyzed in terms of constitutive correctness, even though the fittingness of action can't be. What rationale for this might there be?

One argument for fittingness-as-constitutive-correctness comes from Schroeder (2010). Schroeder observes that both correctness and fittingness are unaffected by WKR's: that WKR's are just as irrelevant to whether a response is constitutively correct as they are to whether a response is fitting. For example, just as the moral or pragmatic benefits of making a chess move are irrelevant to whether the move is correct according to the standard(s) constitutive of chess playing, so too are any benefits

---

<sup>23</sup> See esp. (Lord and Sylvan 2019) and (Sharadin 2015).

<sup>24</sup> Though we disagree with Kiesewetter that RKR's are "essentially linked" to constitutive standards of correctness, we take the present paper to complement his. While Kiesewetter provides primarily positive arguments for thinking that RKR's, and RKR's for belief in particular, are authoritatively normative, our main aim here is to defend this thesis by considering and responding to the most common arguments against it.

<sup>25</sup> See, e.g., Weiner (2005).

<sup>26</sup> Parallel points hold even if the constitutive correctness condition for assertion turns out not to be truth, but instead knowledge (Williamson 2000, Hawthorne 2004) or justification (Kvanvig 2009, 2011).

of admiration irrelevant to the attitude's fittingness. As Schroeder summarizes: "standards of correctness are unaffected by the kinds of incentives that come into play in Wrong Kind of Reasons scenarios" (2010: 33). He conjectures that fittingness just is constitutive correctness.

We have two responses. First, we don't think Schroeder's observation suffices to show that fittingness and correctness are equivalent, much less identical. From the fact that A and B are unaffected by C, it doesn't follow that A and B are the same. My nightstand and clock are equally unaffected by the weather, but they aren't the same object. Similarly, from the fact that fittingness and correctness are equally unaffected by WKR's, it doesn't follow that they're the same relation. Second, Schroeder's conjecture, if true, would equally well support the view that fitting action is equivalent to constitutively correct action. But, as we've just argued, this view seems false.

A second argument that the fittingness of an attitude is a matter of its being constitutively correct may seem more promising. This argument starts with the observation that the conditions under which an attitude is fitting strikingly coincide with the conditions under which it's natural to call the attitude 'correct'.<sup>27</sup> It's natural, for example, to call admiration correct when its object is admirable, to call desire correct when its object is desirable, and to call belief correct when its object is credible.<sup>28</sup> A simple explanation of this is that fittingness just is correctness. Indeed, if fittingness weren't correctness, then it's unclear how the coincidence in question could be plausibly explained.

In response, note first the polysemy of 'correct'. Sometimes 'correct' means accurate or true. Some authors identify correctness in this sense with fittingness, claiming that the fittingness of an attitude is a matter of its accurately representing its object (e.g., Tappolet 2011). This isn't the sense in which a chess move is 'correct' when it satisfies the standards constitutive of chess playing. This latter sense of 'correct' amounts to conformity with a norm, whereas the former sense does not. Hence, there's no sense in which the former sense of 'correct' is normative, whereas there is a clear sense in which the latter is. We reject the view that fittingness is 'correctness' in the sense that amounts to accuracy for reasons well-rehearsed by others (Schroeder 2010, Svavarsdóttir 2014, Naar 2021, D'Arms this volume).<sup>29</sup> So, we're not happy to say that the fittingness of an attitude amounts to its being 'correct' in this sense.

However, we're quite happy to say that fitting attitudes are 'correct' in the second sense, i.e., in the sense that amounts to their satisfying a norm. Indeed, we accept that attitudes are fitting when they satisfy certain norms that govern them. What we deny is that an attitude's being 'correct' in this sense amounts to its satisfying a norm that governs it *constitutively*. For although constitutive norms of correctness are normative in one sense, viz., formally, they aren't in another: they lack authority. So

---

<sup>27</sup> See, e.g., (Schroeder 2012).

<sup>28</sup> It's also natural to call a belief 'correct' when it's true. As we make clear below, in a certain sense of 'correct', this claim is trivial. But as we also make clear, it's not plausible that the kind of correctness to which this sense of the word refers is identical to fittingness; see also the discussion in note 29.

<sup>29</sup> A less discussed reason for resisting this view is that it would make trivial the claim that fitting beliefs are always true beliefs, insofar as it's assumed (plausibly) that beliefs represent their objects as being true. We take this to be a problem for the view since we think it's not at all trivial that fitting beliefs are always true. It's fitting to believe a proposition when it's *credible*, or belief-worthy, and it's a substantive, normative, and controversial matter whether a proposition can be belief-worthy, or merit belief, only if true. Indeed, we find it plausible that a proposition could be credible or belief-worthy even if false—perhaps if one has excellent evidence for it.



it's the thesis that fittingness amounts to *constitutive* correctness, specifically, that threatens its authoritative normativity, and hence that of RKR's. So, it's this thesis we reject. But this is consistent with claiming that when an attitude is fitting, it's correct according to a norm that's *external* to it, i.e., a norm not baked into the attitude's essence. And we're happy to call fitting attitudes 'correct' in this sense. So, our view easily explains the data that it's natural to call fitting attitudes 'correct'.

What's at issue, then, between us and our opponents, is whether the fittingness of an attitude amounts to its being correct according to a norm that governs it internally or externally. In a neglected discussion, Gideon Rosen (2001) remarks on the difficulty of adjudicating the dispute of whether the norms that govern certain mental states are external to them or instead constitutive of their nature. Borrowing from that discussion, one way to think of the issue is this: the view that correctness (or fittingness) norms are internal to the attitudes they govern entails that were God to suspend the correctness norms, the attitudes they govern would also be eradicated; whereas on the view that correctness norms are external to the attitudes, the attitudes would survive. We find the latter idea more attractive, and while we can't offer an extended argument for it here, we note the following two points in its favor: First, the view that we prefer is compatible with a wider range of views about the natures of fit-evaluable attitudes and is in this way more ecumenical. For example, the view that correctness norms are constitutive of attitudes may rule out functionalism or other forms of naturalism in the philosophy of mind, whereas our view is clearly consistent with these positions (Rosen 2001). Second, as noted above, our favored view is compatible with the unity of fittingness as it applies to actions and attitudes whereas our opponent's view is not. As we've already argued, the fittingness of certain actions—e.g., triple axels and assertions—can come apart from their constitutive correctness. (Here, too, we're happy to call these actions 'correct' if they're fitting, but, in light of our arguments, 'correct' here can't mean constitutively correct.) So, again, our opponent is left claiming that the fittingness of actions is a very different property than the fittingness of attitudes with no clear rationale for doing so. Hence, we think that the burden is on our opponent to explain why we're compelled to accept that the fittingness of an attitude amounts to its being internally rather than externally correct, particularly given the disunified view of fittingness that this picture entails.<sup>30</sup>

Our argument here has some important upshots. Because Schroeder and others assume that RKR's arise from constitutive standards of correctness, they typically argue that the distinction between RKR's and WKR's is a general distinction that applies to actions as well as attitudes, by pointing to actions or activities that have constitutive standards—triple axels, tying knots, setting a proper English dining table, etc. This leads Errol Lord and Kurt Sylvan (2019) to claim that there's also a "right kind of reason problem", which is to explain why RKR's for certain attitudes are authoritatively normative while RKR's for action (and perhaps some attitudes) are not. After all, the fact that tying a knot in a

---

<sup>30</sup> In principle, what might motivate the view that norms of fittingness or correctness internally rather than externally govern our attitudes? The best example of a possible motivation that we can find comes from Shah (2003), who argues that the claim that belief is constitutively governed by a norm of truth provides the best explanation of a psychological phenomenon he calls "transparency" (roughly, that in deliberative contexts, only what you take to be evidence can be a motivating reason for belief). Shah (2008) also suggests that there may be a similar motivation for thinking that intention is governed by a constitutive norm. We can't consider Shah's position here, but we find plausible existing criticisms of both its psychological and normative claims (for the former kind of criticism, see esp. McHugh 2015; for the latter, see, inter alia, Steglich-Peterson 2006).

certain way would be a constitutively correct way of tying a slip knot does not seem like a source of authoritatively normative reasons for anyone to tie a knot in the relevant way.

But our argument that RKR's are fit-makers and that fittingness is distinct from constitutive correctness, suggests that actions that have constitutive standards of correctness don't exemplify the distinction between RKR's and WKR's at all. This doesn't imply that the distinction isn't a general one that applies to action too—we think it does. But it does imply that in order to find actions that exemplify the distinction, we need to look for actions that seem *merited* as a response, regardless of whether performing the action would be of any value.<sup>31</sup> Relatedly, our argument suggests that there's no right kind of reason problem. If constitutive-correctness-makers for action aren't RKR's, as we claim, then they don't present a challenge to explain why RKR's for attitudes are authoritatively normative while these merely formally normative reasons for action are not. On our view, *all* RKR's are authoritatively normative reasons because they're fit-makers and not constitutive-correctness-makers.

### 3. Formal features of right-kind reasons

Another theoretical argument against the authoritative normativity of RKR's has its roots in a recent paper by Barry Maguire (2018). This argument targets the authoritative normativity of RKR's for affective attitudes in particular. Maguire argues that authoritative normative reasons are essentially gradable and contributory. They're contributory in that they constitute “incomplete parts of a specific kind of explanation of overall normative facts, such as facts about what you ought to do” (2018: 780). In particular, reasons contribute to “weighing explanations” of overall normative facts by competing and combining with one another. Reasons are gradable in the sense that they have weights, such that one reason might be weightier than another or provide a greater degree of support for the response it favors. But according to Maguire, fit-making facts for affective attitudes are neither gradable nor contributory: they don't have weights and they never contribute to explanations of overall normative facts by competing and combining. Maguire thus concludes that fit-making facts, or RKR's, for affective attitudes aren't really reasons. So, assuming something can have normative authority only if it is or provides a reason, it follows that RKR's for affective attitudes aren't authoritatively normative.

In principle, there are several ways to reply to this argument. One strategy would be to deny that something can have normative authority only if it is or provides a reason.<sup>32</sup> Alternatively, one might reject that reasons are essentially gradable and contributory, or instead argue that RKR's for affective attitudes do have these features. In this section, we defend the authoritative normativity of RKR's by pursuing this latter option: we grant that something can have normative authority only if it is or provides a reason, and that reasons are essentially gradable and contributory, but argue that fit-making facts for affective attitudes also possess these features essentially. We aim not only to diagnose where Maguire goes wrong in his arguments that these facts lack these features, but to suggest a positive view about what determines the weights of fit-making facts for affective attitudes and how such facts compete and combine with one another so as to contribute to weighing explanations of overall normative facts.<sup>33</sup>

---

<sup>31</sup> For some discussion, see (Howard 2021).

<sup>32</sup> This possible response was suggested to us by Jonathan Way.

<sup>33</sup> For other recent responses to Maguire's arguments, see esp. (Faraci 2020) whose responses we largely agree with and build upon here, and (Heape 2020), which provides a fruitful exploration of the prospects for reasons-firsters of adequately responding to Maguire's arguments; see also (McHugh and Way *fc*, ch. 7)

Let's start with Maguire's argument that fit-making facts never compete, which proceeds by cases. First, Maguire considers the loss of his grandmother. Upon her passing, he felt both sadness and relief. The fact that his grandmother was suffering seemed to make the latter response fitting; the fact that she was the family matriarch and "had great chat" up until the end seemed to make the former fitting (ibid.). Intuitively, these considerations don't compete with each other. As Maguire puts it: "The facts that made the sadness fitting didn't make the relief unfitting, nor did the facts that made the relief fitting make the sadness unfitting" (ibid). It was fitting to feel both sadness and relief—sadness that the family matriarch had passed and relief that she was no longer suffering.

Second, Maguire considers a case in which your friend Andrew gets a promotion that you're also up for. According to Maguire, the fact that your friend got the promotion makes it fitting to feel pleased he got it and the fact that you didn't get the promotion makes it fitting to feel disappointed you didn't. Again, Maguire claims, "the considerations supporting these attitudes do not compete. The fact that you didn't get the promotion doesn't make it unfitting to feel pleased and the fact that your friend got it doesn't make it unfitting to feel disappointed. It is fitting to feel pleased *and* disappointed in these different respects in this case" (2018: 787-88). Maguire concludes that fit-making facts never compete, but instead directly, by themselves, each make a specific and separate attitude fitting.

We agree with Maguire that, in the cases he considers, there's no competition among the fit-making facts. But we deny that these cases show that fit-making facts for affective attitudes never compete and that these facts aren't reasons. As David Faraci (2020) observes, Maguire's argument seems to presuppose that reasons compete *per se*. But they don't. Reasons for me to get work done today don't compete with the reasons for me to relax tomorrow; and reasons for me to believe it'll rain this afternoon don't compete with reasons for me to hope it won't. Instead, reasons compete only insofar as they favor incompatible alternatives. Faraci suggests that the explanation for why the fit-making facts in Maguire's cases don't compete is that the attitudes in these cases aren't incompatible alternatives. We agree. Feeling sad that the family matriarch has passed isn't incompatible with feeling relief that she's no longer suffering. Likewise, feeling disappointed that you didn't get the promotion isn't incompatible with feeling pleased that your friend did. So, the lack of competition between the fit-making facts in Maguire's cases doesn't suggest that these facts aren't reasons for the attitudes they make fitting.<sup>34</sup>

What Maguire needs to show is that fit-making facts for affective attitudes don't compete even when the attitudes they make fitting are alternatives. But we think this can't be shown. We'll suggest that fit-making facts do compete in such cases and we'll explain how this works in practice.

First, what does it take for two or more attitudes to be incompatible alternatives? In the case of action, two acts are incompatible alternatives when the performance of one makes impossible the performance of the other. In the case of attitudes, things are different. It's possible both to believe *p* and to disbelieve *p*, but these attitudes are, in some sense, alternatives.<sup>35</sup> Our working hypothesis is that affective attitudes are alternatives to one another when they can't simultaneously be fitting.<sup>36</sup>

---

<sup>34</sup> McHugh and Way (fc, ch. 7) also make this point.

<sup>35</sup> Selim Berker makes this point in an earlier draft of (Berker 2018).

<sup>36</sup> Faraci (2020: 228) suggests a similar account of when two attitudes are incompatible alternatives. This account looks plausible in the case of affective attitudes and, we think, belief, but less so when it comes to certain conative attitudes such as intention. For example, it can at once be fitting to intend to *A* and to intend

For each type of attitude that can be fitting, there's some evaluative property such that the fittingness of the attitude is equivalent to its object possessing that property. For example, admiration is fitting just in case its object is *admirable*, desire is fitting just in case its object is *desirable*, and awe is fitting just in case its object is *awesome*. Given this, we have a nice test for whether two attitudes can't simultaneously be fitting and hence whether they're in this way alternatives: Two attitudes can't simultaneously be fitting if it's impossible for their object(s) to bear simultaneously the evaluative properties to which the fittingness of each is equivalent. For example, a person can't simultaneously be both admirable and deplorable, and so it can't at once be fitting both to admire and to deplore them. Hence, admiring and deploing one and the same thing are alternatives in the relevant sense.

Of course, a person might have a mix of admirable and deplorable properties. If they do, then it can simultaneously be fitting to admire the admirable properties and to deplore the deplorable ones: the deplorability of one of a person's properties needn't preclude the admirability of some other property she possesses. So it can be fitting, for example, to admire Angie's brilliance while deploing her evil intentions. But it can't at once be fitting both to admire and deplore Angie herself.

Could it be fitting both to admire and deplore one and the same property of a person? No. Angie's brilliance can't be admirable *and* deplorable. Plausibly, if Angie is brilliant, her brilliance is admirable; hence, not deplorable. But what if Angie is brilliant, but puts her brilliance to bad use? Then, we think, it's fitting to admire her brilliance *as such*, but to deplore the use to which she puts it.

We're now in a position to give a deeper explanation of why there's no competition among the fit-making facts in Maguire's cases, viz., that the affective attitudes his cases involve can simultaneously be fitting and so don't count as alternatives. For example, the disappointingness of the fact that you didn't get the promotion is compatible with the joyousness of the distinct fact that your friend did. Hence, disappointment regarding the former fact and joy regarding the latter fact don't count as alternatives. Since there can be competition among fit-making facts for attitudes only if the attitudes are alternatives, this is why there's no competition among the fit-making facts in Maguire's cases.

So, how does competition among fit-making facts work when the relevant attitudes are alternatives? Suppose Angie has a mix of admirable and deplorable properties. What's the fitting attitude toward Angie herself? In particular, is it fitting to admire Angie, to deplore her, or to feel a kind of ambivalence? We think the answer depends on the outcome of a weighing process that takes as inputs Angie's admirable and deplorable features. Here's a brief sketch of how this works. Each of Angie's admirable features is admirable to a degree. The degree to which an admirable feature is admirable corresponds to its weight. The same goes for Angie's deplorable features: each is deplorable to a degree, and the degree to which a deplorable feature is deplorable equals its weight. To determine whether Angie herself is admirable or deplorable or neither, we first sum the weights of her admirable features, and then sum the weights of her deplorable features. We then determine the difference between these values. If the difference meets some (possibly vague) threshold, such that the combined weight of Angie's admirable features is sufficiently greater than the combined weight of her deplorable

---

to not-A (when Aing and not-Aing are each worth doing) but these attitudes seem clearly to be alternatives to each other. Perhaps a more unified account of when attitudes are alternatives is forthcoming, but since our focus (and Maguire's) is affective attitudes in particular, the above hypothesis about when such attitudes are alternatives will suffice for present purposes. Thanks to Selim Berker, Jonathan Way, and Alex Worsnip for helpful conversation and correspondence about these issues.

ones, then Angie herself is admirable, and so fitting to admire.<sup>37</sup> If, on the other hand, the combined weight of Angie's deplorable features sufficiently exceeds the combined weight of her admirable ones, then she is deplorable, and so fitting to deplore. And, finally, if neither of these conditions is met, then Angie is neither admirable nor deplorable, and so it's not fitting either to admire or to deplore her. Instead, in this last kind of case, it's fitting to be ambivalent towards Angie, where ambivalence is something like the affective analogue of suspending judgment about a proposition's truth.

Several points of clarification are in order. First, settling whether Angie is admirable (or deplorable) doesn't settle *how* admirable she is. There are many views about how to determine how admirable Angie is that are compatible with the above account, but we tend to favor this one: the extent to which Angie is admirable is determined by how much she exceeds the threshold for being admirable (rather than deplorable). Roughly, once Angie has met the threshold for being admirable, the greater the degree to which she exceeds that threshold, the more admirable Angie is.

Second, one might wonder how our model works in cases where an affective attitude has several alternatives.<sup>38</sup> For instance, just as someone can't be both admirable and deplorable, it's also plausible that they can't be both admirable and despicable. So, just as deploring someone is an alternative to admiring them, so too is despising them. We suggest that someone is admirable, rather than deplorable or despicable, if the combined weight of their admirable features is sufficiently greater than each of the combined weights of their deplorable features and their despicable features. And if someone is admirable, they're fitting to admire, rather than to despise or deplore. In this way, our model can deliver verdicts about which attitudes are fitting, even in cases involving attitudes with multiple alternatives.<sup>39</sup>

Our model thus explains how facts about features that can contribute to the fittingness of alternative affective attitudes can compete to determine which attitude is fitting. In other words: it explains how fit-making facts, or RKRs, for alternative affective attitudes can compete to determine the (un)fittingness of each. We think this model applies not only in cases like the above, where the common object of a set of alternative attitudes is a person (or some other concrete object), but also in cases where the common object of a set of alternatives is a complex state of affairs, event, process, etc. For example, recall Maguire's case in which your friend Andrew gets a promotion that you were also considered for. Let the fact that you didn't get the promotion be a complex state of affairs that includes the fact that you won't get the raise but also the fact that you won't work for longer hours. The fact that you won't get the raise is disappointing, but the fact that you won't work for longer hours is pleasing. So, on our model, whether you not getting the promotion is disappointing, and thus whether it's fitting for you to

---

<sup>37</sup> Why "sufficiently greater" rather than just "greater"? To preclude the possibility that Angie herself could be admirable in a case where her features are just slightly more admirable than they are deplorable.

<sup>38</sup> Thanks to Selim Berker for raising this question.

<sup>39</sup> In cases like this, we think it's most plausible to hold that how admirable someone is a function both of how much they exceed the threshold for being admirable rather than deplorable and how much they exceed the threshold for being admirable rather than despicable. We're tempted by the view that in such cases no precise degree of admirability can be specified, and instead the best we can do is to specify an imprecise range.

be disappointed by this complex state of affairs, is determined by whether its proper parts are sufficiently more disappointing than they are pleasing. It's easy to see how this picture might generalize.<sup>40</sup>

This model of how fit-making facts can compete also reveals how such facts can combine, such that they sometimes together contribute to explanations of an attitude's fit. We therefore conclude, contrary to Maguire, that fit-making facts, or RKR, for affective attitudes are contributory: they contribute to weighing explanations of overall normative facts by competing and combining.

Our model also substantiates the fact that fit-making facts for affective attitudes are gradable, i.e., that they have weights. The weight of a fit-making fact for an attitude depends on the extent to which the fact can contribute to making it the case that the object of the attitude has the evaluative property to which the fittingness of the attitude corresponds. For example, loving your daughter is fitting just in case she's lovable. On our model, your daughter is lovable just in case the combined lovable of her lovable features is sufficiently greater than the combined unlovability of her unlovable ones. So, suppose your daughter's quirky sense of humor is more lovable than her intellectual curiosity. Then, all else equal, her quirky sense of humor can play a greater role than her intellectual curiosity in making it the case that she's lovable. Hence, the former is a stronger RKR to love her.

Why does Maguire think that fit-making facts aren't gradable and where does his argument go wrong? Maguire claims that fit-making facts aren't gradable because, strictly speaking, an attitude can't be more or less fitting. For example, in comparing the deaths of an octogenarian and a twenty-year old, Maguire claims that while we might say that the tragedy of the youngster's death makes it more fitting to feel sad about her death, this just amounts to claiming that the tragedy of the youngster's death makes it fitting to feel more sadness (2018: 790). And although the fittingness conditions for attitudes are themselves gradable—something can be more or less admirable, fearsome, etc.—Maguire insists that this too only makes it fitting to have certain attitudes to varying degrees: if  $x$  is more admirable or fearsome than  $y$ , then it's fitting to admire or fear  $x$  more than  $y$ . So, Maguire insists that fittingness is an all or nothing normative property and he takes this to show that RKR aren't gradable.

The problem with Maguire's argument is that fittingness needn't be gradable in order for fit-making facts to be gradable. To see the point, consider reasons to act. To say that R1 provides a stronger reason to perform act A than R2 is to say that R1 (non-factively) contributes to making it the case that one ought to or may A to a greater extent than R2. But notice that the overall normative statuses of which normative reasons provide contributory explanations—what you ought to or may do—are all or nothing. Either you ought to (or may) do A or not.<sup>41</sup> The same may hold for fit-making facts and fittingness: fittingness may be an overall, all or nothing normative status, even though fit-making facts are gradable in the sense that one fact can contribute to making it fitting to have some attitude more so than some other fact does. For example, as explained above, your daughter's quirky sense of humor might contribute to making it fitting to love her more so than her intellectual curiosity.

So, even if fittingness isn't gradable, this doesn't entail that fit-making facts aren't gradable. And even if the fitting intensity of an attitude corresponds to the degree of the relevant evaluative

---

<sup>40</sup> Our model leaves open whether fit-making facts for affective attitudes with different objects could ever compete. Whether this is possible depends on whether such attitudes can ever count as alternatives. If they can, then our model can also be generalized easily to explain how the fit-making facts for such attitudes compete.

<sup>41</sup> Faraci (2020: 232) makes this point with respect to permissibility.

property that the attitude's object possesses, this also doesn't suggest that fit-making facts for the attitude aren't gradable. This is because, as our model explains, this too is compatible with taking the particular features of an attitude's object to be gradable fit-makers that determine not only the fitting intensity of the attitude in question, but also whether that attitude is fitting in the first place. For example, the particular (un)lovable features of your daughter determine not only how much love it's fitting for you to feel toward her, but also whether your daughter is lovable, or fitting to love, in the first place.

#### 4. Conclusion

We've argued that the main reason for thinking that RKR's are authoritatively normative is simply that it's required to explain the intuitive data. Rejecting this data is a serious cost that should be embraced only if there are strong theoretical reasons for thinking that, despite the appearances, RKR's aren't authoritatively normative after all. And we've argued that the main theoretical reasons for thinking this that have been offered in the literature so far aren't compelling: RKR's needn't and shouldn't be analyzed in terms of constitutive correctness and they have the essential features of normative reasons. We therefore conclude that skepticism about the authoritative normativity of RKR's is unwarranted.

#### References

- Côté-Bouchard, C., & Littlejohn, C. (2018). Knowledge, reasons, and errors about error theory. In C. Kyriacou & R. McKenna (eds.), *Metaepistemology, Realism, and Anti-Realism*, 147-171. Palgrave Macmillan.
- D'Arms, J. (fc) Fitting emotions. In C. Howard & R. Rowland (eds.), *Fittingness*. Oxford: Oxford University Press.
- D'Arms, J., & Jacobsen, D. (2006). Anthropocentric constraints on human value. *Oxford Studies in Metaethics*, 1, 99-126.
- Berker, S. (2018). A combinatorial argument against practical reasons for belief. *Analytic Philosophy*, 59, 427-470.
- Berker, S. (fc) The deontic, the evaluative, and the fitting. In C. Howard & R. Rowland (eds.), *Fittingness*. Oxford: Oxford University Press.
- Brandt, R. (1946). Moral valuation. *Ethics*, 56, 106-121.
- Brentano, F. (1889/2009). *The origin of our knowledge of right and wrong*. Oxford: Routledge.
- Broad, C.D. (1930). *Five types of ethical theory*. London: Routledge.
- Cherry, M. (2018). The errors and limitations of our "anger-evaluating" ways. In *The Moral Psychology of Anger*, edited by Myisha Cherry and Owen Flanagan, 49-66, Lanham: Rowman and Littlefield.
- Ewing, A.C. (1948). *The definition of good*. London: Routledge & Kegan Paul.
- Faraci, D. (2020). We have no reason to think there are no reasons for affective attitudes. *Mind*, 129, 225-234.
- Feinberg, J. (1970). *Doing & deserving: Essays in the theory of responsibility*. Princeton, NJ: Princeton University Press.
- Gibbard, A. (1990). *Wise choices, apt feelings*. Cambridge, MA: Harvard University Press.
- Gertken, J., & Kiesewetter, B. (2017). The right and wrong kind of reasons. *Philosophy Compass*, 12, e12412.

- Hawthorne, J. (2004). *Knowledge and lotteries*. Oxford: Oxford University Press.
- Heape, A. (2020). How there could be reasons for affective attitudes. *Ethical Theory and Moral Practice*, 23, 667-680.
- Howard, C. (2016). In defense of the wrong kind of reason. *Thought*, 5, 53-62.
- Howard, C. (2018). Fittingness. *Philosophy Compass*, 13, e12542.
- Howard, C. (2019a). The fundamentality of fit. *Oxford Studies in Metaethics*, 14, 216-236.
- Howard, C. (2019b). Weighing epistemic and practical reasons for belief. *Philosophical Studies*, 177, 2227-2243.
- Howard, C. (fc). Consequentialists must kill. *Ethics*.
- Kelly, T. (2003). Epistemic rationality as instrumental rationality: A critique. *Philosophy and Phenomenological Research*, 66, 612-640.
- Kiesewetter, B. (fc). Are epistemic reasons normative? *Nous*.
- Kvanvig, J. (2009). Assertion, knowledge and lotteries. In P. Greenough & D. Pritchard (eds.), *Williamson on knowledge*, 140-60. Oxford: Oxford University Press.
- Kvanvig, J. (2011). Norms of assertion. In J. Brown & H. Cappelen (eds.), *Assertion: New philosophical essays*, 233-250. Oxford: Oxford University Press.
- Leary, S. (2017). In defense of practical reasons for belief. *Australasian Journal of Philosophy*, 95, 529-42.
- Lord, E., & Sylvan, K. (2019). Reasons: Wrong, right, normative, fundamental. *Journal of Ethics and Social Philosophy*, 15, 43-74.
- Maguire, B. (2018). There are no reasons for affective attitudes. *Mind*, 127, 779-805.
- Maguire, B., & Woods, J. (2020). The game of belief. *The Philosophical Review*, 129, 211-49.
- Mantel, S. (2019). Do epistemic reasons bear on the ought simpliciter? *Philosophical Issues*, 29, 214-27.
- McDowell, J. (1998). *Values and secondary qualities*. Reprinted in his *Mind, Value and Reality*. Cambridge, MA: Harvard University Press.
- McHugh, C. (2015). The illusion of exclusivity. *European Journal of Philosophy*, 23, 1117-36.
- McHugh, C., & Way, J. (2016). Fittingness first. *Ethics*, 126, 575-606.
- McPherson, T. (2018). Authoritatively normative concepts. *Oxford studies in metaethics*, vol. 13.
- Papineau, D. (2013). There are no norms of belief. In T. Chan (ed.), *The Aim of Belief*, 64-79. Oxford: Oxford University Press.
- Parfit, D. (2011). *On what matters*, vol. 1. Oxford: Oxford University Press.
- Reisner, A. (2008). Weighing pragmatic and epistemic reasons for belief. *Philosophical Studies*, 138, 17-27.
- Reisner, A. (ms). The pragmatic foundations of theoretical reason.
- Rinard, S. (2019). Equal treatment for belief. *Philosophical Studies*, 176, 1923-1950.
- Rosen, G. (2001). Brandom on modality, normativity, and intentionality. *Philosophy and Phenomenological Research*, 63, 611-623.
- Rowland, R. (fc). The authoritative normativity of fitting attitudes. *Oxford Studies in Metaethics*.
- Schroeder, M. (2007). *Slaves of the passions*. Oxford: Oxford University Press.
- Schroeder, M. (2010). Value and the right kind of reason. *Oxford Studies in Metaethics*, vol. 5, 25-55.
- Shah, N. (2003). How Truth Governs Belief. *The Philosophical Review*, 112, 447-482.
- Sharadin, N. (2015). Reasons wrong and right. *Pacific Philosophical Quarterly*, 96, 1-29.
- Sidgwick, H. (1890). *The methods of ethics* (4<sup>th</sup> ed.). London: Macmillan.



- Skorupski, J. (2010). *The domain of reasons*. Oxford: Oxford University Press.
- Srinivasan, A. (2018). The aptness of anger. *Journal of Political Philosophy*, 26, 123-144.
- Steglich-Peterson, A. (2006). No norm needed: On the aim of belief. *Philosophical Quarterly*, 56, 499-516.
- Velleman, D. (1992). The Guise of the Good. *Nous*, 26, 3-26.
- Way, J. (2012). Transmission and the wrong kind of reason. *Ethics*, 122, 489-515.
- Way, J. (2013). Value and reasons to favour. *Oxford Studies in Metaethics*, 8, 27-49.
- Weiner, M. (2005). Must we know what we say? *The Philosophical Review*, 114, 227-51.
- Whiting, D. (2021). Aesthetic reasons and the demands they (do not) make. *Philosophical Quarterly*, 71, 407-427.
- Wiggins, D. (1987). *A sensible subjectivism?* In his *Needs, Values, Truth: Essays in the Philosophy of Value*. Oxford: Blackwell.
- Williamson, T. (2000). *Knowledge and its limits*. Cambridge, MA: Harvard University Press.
- Wodak, D. (2019). Mere formalities: normative fictions and normative authority. *Canadian Journal of Philosophy*, 49, 828-850.
- Wodak, D. (fc). Approving on the basis of normative testimony. *Oxford Studies in Metaethics*.
- Woods, J. (2018). The authority of formality. *Oxford Studies in Metaethics*, vol. 13, 217-229.