

Normativity as reactive shield

Singer's 'Famine, Affluence, and Morality' (1972) makes students wonder why they should be moral. Glaucon wondered too (*Republic* 361a-c). But the question is strangely elusive. They're not asking if, given the demands of morality, they should be moral. That question is too trivial to be their target. Nor are they asking whether they should be moral given the demands of *prudence*, as H. A. Prichard (1912) rightly emphasized. In any case, they might similarly wonder why they should be prudent.

Christine Korsgaard (1996) called the question 'why should I be moral?' the *normative* question. We make affirmative judgments with the same kind of normative 'should': "you are in some *Burning Hotel*, and you can save your life only by jumping into some canal. I am outside your hotel, which I know to be on fire, and I can see you at some window above the canal" (Parfit 2011, 283). Parfit thinks that you should jump. He takes the thought that you should jump to be a *normative* thought, a thought with the normative 'should'.

A powerful and influential account of what is it to be normative centers motivation:

the PRACTICAL CONSTRUAL: the fact that I should ϕ according to some standard is normative iff I'd be motivated to ϕ if I'm rational and accept that fact.

This formulation assumes that there are facts about what I should do relative to different standards – I should keep my promises, given the standard of morality, and I should address certain people as 'Sir', given the demands of feudal honor. It distinguishes some of those 'should' facts as normative by their link to rational motivation. Rational people feel motivated to keep their promises, but not necessarily to address certain people as 'Sir'. Kantians like Korsgaard, expressivists like Allan Gibbard (1990), and Humeans like Bernard Williams (1979) all accept versions of the Practical Construal. The Construal makes good sense of the question 'why be moral?'; it's a question about what we're motivated to do if rational. Parfit joins Stephen Darwall in insisting that moral naturalists *must* accept the

Practical Construal: “ ‘for the philosophical naturalist, concerned to place normativity within the natural order, there is nothing plausible for normative force to be other than motivational force’ ” (Parfit 2011, 363). If they’re right, reductive naturalism about *normativity* must center psychological states like desires; the prospects for reductive naturalism about normativity threaten to hang on the prospects for some kind of broadly Humean view.¹

This paper develops an alternative. The alternative descends from Mill’s classic account of moral wrongs as what ought to be sanctioned. Mill’s classic account looks circular if repurposed as a general account of normativity, because it seems forced to construe the *ought normatively*. I reformulate Mill’s account to avoid objectionable circularity. I suggest, roughly, that a fact’s normativity consists in its constitutively governing certain sanctions. I aim to show that the resulting construal of normativity is on a par with the Practical Construal. Establishing parity helps us ask which metaethical questions should be central. In trying to establish parity, I’m trying to identify functional roles that might distinguish normativity. I won’t ask anything deeper, like if one functional role uniquely *constitutes* normativity.²

1 Normativity, what

This paper aims to identify what Korsgaard calls the normative question. Identifying the normative question doesn’t mean answering it. Korsgaard and Gibbard identify the normative question similarly but answer it differently.

I’m looking for an alternative to the Practical Construal because the Construal threatens a central realist commitment: that moral facts are independent of and more fundamental than our individual attitudes. The Construal requires a kind of ‘motivational internalism’ that links normative facts to motivation. Now many realists do attempt to vindicate motivational internalism. But their attempts end up costly – costly enough to throw the realist commitment into doubt. I suspect that realists would do better to reject motivational internalism. Despite that suspicion, I’ve found myself blithely appealing to the Practical Construal when teaching, when students

¹The alternatives seem to reject the distinction between what’s normative and what’s not (Baker 2018; Copp 1997; Foot 1972; Sagdahl 2014; Tiffany 2007).

²As a result, this paper is compatible with the more general skeptical worry that metanormative theory might not share any ultimate common subject matter, as Stephen Finlay (2019) has recently discussed.

challenged me to explain the difference between morality and a society's moral conventions. In the moment, the Practical Construal felt inescapable; it felt like the only real way of explaining the difference.

In particular, the Practical Construal felt like the only way of satisfying three central desiderata. The first desideratum is that a construal of normativity interprets Korsgaard's normative question as a *substantive* question. I opened with this desideratum; I emphasized that students who ask 'why be moral?' aren't asking something trivial. We can sharpen this question by contrasting moral norms with feudal norms of address and respect. We don't think the feudal norms are genuine norms but allow that moral norms could be genuine norms. Korsgaard's normative question could then be formulated as concerning the difference between moral norms and feudal norms (Dorsey 2016, 9).

The second desideratum, which is closely related, is that the identification illuminates when a group is using a word or concept normatively. Horgan and Timmons give their 'moral twin earth' objections by imagining groups using words normatively, imagining a group of Martians "being much like humans in their cognitive sophistication and their social institutions, Martians too employ moral terms and concepts; their moral vocabulary is intertranslatable with our own" (Horgan and Timmons 1996, 21). They object that various kinds of moral realism have implausible upshots about that race of Martians. David Copp (2000) develops one kind of response to them. Horgan and Timmons interpret his answer as claiming that *what it is* to use 'wrong' normatively determines the same unique referent for any community's use of a term translatable as 'wrong'. Adequate construals of the normative question would illuminate this dispute.

The third desideratum is that the identification explains our intuitive grasp of what's normative. The demands of prudence and morality strike most of us as genuinely normative; epistemic norms do too. In contrast, the demands of feudal honor and of etiquette don't seem genuinely normative. Construals of the normative question become more plausible by vindicating or at least explaining these convictions. In addition, some but not all of us see some legal and aesthetic norms as genuinely normative. Construals of the normative question become more plausible by illuminating why.

The Practical Construal secures all three desiderata.

the PRACTICAL CONSTRUAL: the fact that I should ϕ according to some standard is normative iff I'd be motivated to ϕ if I'm rational and accept that fact.

It explains how the normative question is a substantive question:

it's about rational motivation to be moral. It also supports an elegant account of when someone uses a concept normatively: they use it normatively when their rational use of it motivates them. In fact, Eklund only considers versions of the Practical Construal when imagining groups using different normative concepts – considering, for instance, the possibility that “a predicate is normative because it can be conventionally used to guide action” (Eklund 2020, 130). The Practical Construal can seem *inescapable*.

The Practical Construal also vindicates our intuitive grasp of what's normative. A rational person would plausibly be motivated to ϕ , at least to some extent, if she judged that prudence enjoined ϕ -ing – similarly for a rational person who judges that morality enjoined ϕ -ing. Now morality might forbid something that prudence demands, and morality might trump. The Practical Construal can interpret morality's trumping here as consisting in the fact that a rational person would in this case be motivated by morality rather than prudence. Though epistemic norms are slightly more complicated, similar points hold, as Scanlon notes: “when a rational creature judges that the reasons she is aware of count decisively against a certain attitude, she generally does not have that attitude, or ceases to have it if she did so before—ceases to feel conviction in regard to the belief or to use it as a premise” (Scanlon 1998, 24). Scanlon is focusing here on what an agent *normatively* ought to believe, in talking about reasons that count decisively against the attitude.

Now the Practical Construal *isn't* an all-purpose recipe for classifying an arbitrary fact as normative or not – for instance, it doesn't itself say anything about reasons. It instead simply says when an *ought* fact is normative.³ My construal of normativity also only sorts *ought* facts into the normative bucket or the non-normative bucket. It won't say when reasons are normative, either. I won't explore how to generalize these construals of normativity beyond *ought* facts, because both construals have the same range of options. Maybe, for instance, reasons are normative when they support or explain or ... normative *oughts*.

Proponents of the Practical Construal often find its verdicts about prudence, morality, and epistemology independently plausible. Even if we didn't accept the Practical Construal, we'd independently think that judgments about those domains can motivate rational agents.

The Practical Construal also explains why etiquette and the norms of feudal honor are *not* normative. Etiquette and feudal honor can

³The passage from Scanlon focuses on those sorts of facts, too, since decisive reasons against some attitude mean that you ought not have the attitude.

fail to motivate me without impugning my rationality; I'm not *irrational* if they leave me cold. Again, this point reflects a conviction that's separate from the Practical Construal. Given the demands of feudal honor, I ought to address Ralph Wedgwood as 'Sir'. I can perhaps imagine someone who insists that failures to address him as 'Sir' are irrational. But I'm confident that such a person would be mistaken. Given that confidence, the Practical Construal predicts that that demand of feudal honor isn't normative. It also promises to illuminate whether legal norms or aesthetic norms are genuinely normative, by asking if they could fail to motivate a rational person.

Others have attempted to characterize when standards are normative. One such account takes standards to be genuinely normative if they're *categorical* – if they apply whatever the agent's ends. Philippa Foot (1972) eliminated that account by pointing out etiquette could be categorical.⁴

Others are skeptical about normativity as such. Those skeptics reject the distinction between standards that are genuinely normative and standards that aren't. David Copp (1997, 2020) gave an influential case for skepticism that Derek Baker (2018), Matthias Sagdahl (2014), and Evan Tiffany (2007) bolster. Vocabulary differs. I'm using 'normative', 'genuinely normative', and 'normativity'; I'm interested in characterizing normative obligations. I might also talk about *authoritative* obligations, or perhaps *all-things-considered* obligations, or, more datedly, *rational* obligations.

2 Introducing the Reactive Construal

The Practical Construal generates questions that structure vast swathes of metaethics. For one thing, it requires motivational internalism, which makes trouble for moral realists. And it pulls us back to motivational internalism even in the face of counterexamples to internalism, like those that Sigrun Svavarsdottir (1999) develops. Simon Blackburn suggests that “that externalists can win individual battles. ... But internalists win the war for all that, in the sense that these cases are necessarily parasitic, and what they are parasitic upon is a background connection between ethics and motivation” (Blackburn 1998, 61). Internalists win the war if we accept the Practical Construal. The challenge I'll take up is: how might we explain distinctively normative *oughts* to our students without it? So I want

⁴Stephen Finlay (2019) builds on Joyce's work in generalizing this objection to Parfit's suggestion that genuine normativity is reason-involving rather than rule-involving (Parfit 2011, 144-5).

an alternative to the Practical Construal because I want to open another front in the war on internalism. And I hope that this new front gives us a shot at winning the whole war. That's not the only reason to care about my alternative – but I'm following Enoch (2011) in saying my real motivation.

My official aim is more modest. I just want to convince you that my alternative is a genuine alternative; it's on a par with the Practical Construal. Appreciating parity lets us see how we might escape the Practical Construal.

My alternative builds from Mill's suggestion that

we do not call anything wrong, unless we mean to imply that a person ought to be punished in some way or other for doing it— if not by law, by the opinion of his fellow-creatures; if not by opinion, by the reproaches of his own conscience. (Mill 1863, 48-9)

The paper will focus on sanctions like resentment, indignation, and so on – what Strawson called the *reactive* attitudes. It distinguishes normative facts by the way they govern those reactive sanctions. Because the construal focuses on the reactive attitudes, I'll call it the *Reactive* Construal.

Construals that build from the ideas from Mill and Strawson face at least two important problems. The first problem is that Mill-Strawsonian construals of normativity threaten to be objectionably circular. For instance, Darwall defends a Millian account of *moral* wrong in particular:

for an action to be morally wrong is for it to warrant blame, should the agent lack an adequate excuse (Darwall 2017, 5)

Kieran Setiya (2021) argues in detail that Darwall's suggestion is objectionably circular. When actions warrant blame in the way that interests Darwall, they warrant blame *because* they're morally wrong. The danger of circularity deepens when we generalize this kind of account into an account of normativity in general. Surely it's at least *normatively* warranted blame that matters!

This paper avoids circularity by developing a construal of normativity that rests on *constitutive* norms on reactive attitudes. Constitutive appropriateness isn't normative appropriateness, as Judith Jarvis Thomson emphasizes: “the rules of a game impose no obligation of any kind on the players. The rules of chess do not tell you

what you are under an obligation to do. They do not tell you what you must, or even what you ought to do. Suppose you are playing chess, and it is your turn to move. You then learn that if you don't move your bishop horizontally, hundreds will die! Are you all the same under an obligation to not do so? Must you, ought you not do so? That idea is just silly" (Thomson 2008, 90).

2.1 Identifying the constitutive norms

Much of the paper will develop my official candidate for a constitutive norm on reactive attitudes. But before I introduce my official candidate, I'll reformulate Darwall's approach to illustrate how it could incorporate constitutive norms rather than genuinely normative norms. That reformulation would hold that there is a class of facts with the form 'x ought to ϕ ' that:

- can make blaming x for ϕ -ing constitutively appropriate, absent an excuse.
- can make guilt over x's ϕ -ing constitutively appropriate, absent an excuse.
- ...

Then we can give a Darwall-Style construal of normativity, where that class of facts just is the class of normative facts.

DARWALL-STYLE CONSTRUAL the fact that x ought to ϕ according to some standard is normative iff it makes [blame/ guilt/ ...] for x's failing to ϕ *constitutively* appropriate, absent an excuse.

This construal might avoid the first problem for Millian accounts, of offering a circular characterization of normativity. The righthand side appeals only to constitutive facts about reactive sanctions like blame. (But maybe not: isn't 'excuse' a normative notion?)

The second kind of problem for views like Darwall's is that they seem to *overmoralize* the reactive attitudes. Dale Dorsey illustrates this kind of problem an example from the baseball movie *Bull Durham*:

Crash Davis (played by Kevin Costner, with searing indignation) is chastising Nuke LaLoosh (played by Tim Robbins) for failing to live up to his talent and promise. Importantly, it's clear that he is blaming him not on the

grounds that he has failed himself—Davis holds that this is explicitly ‘your problem’. But the problem he has, his reason for blaming LaLoosh is that LaLoosh doesn’t respect the game. ... And as a matter of pure, first-order intuition, it seems to me that Davis’ blame—assuming his account of LaLoosh and his behavior is right—is perfectly appropriate, ‘fitting’ (Dorsey 2020, 687)

If Davis’ blame is appropriate, the Darwall-Style Construal would predict that the norms of baseball excellence are genuinely normative. And that prediction looks wrong – baseball’s just a game! The Darwall-Style Construal faces this problem because it specifies *whenever* blame is appropriate. Since Davis’ blame looks appropriate, it must arise from genuinely normative standards. But in this case, appropriate blame seems to arise from somewhere more prosaic: Davis’ own commitment to the standards of baseball. I’ll say the Darwall-Style Construal is *excessively moralized*, because ‘excessively normativized’ is too ugly.⁵

This paper tries to avoid overmoralizing by identifying a modest common core of Millian construals. The modest common core only specifies one way for blame to be **in**appropriate. Since it’s only giving conditions for inappropriate blame, it’s simply silent about Dorsey’s verdicts. The disagreement between Darwall and those who find his kind of account excessively moralized is simply *irrelevant* my Millian construals of normativity. Given my candidate, accepting a Millian construal of normativity doesn’t require settling the disagreement in Darwall’s favor.

To warm up, note that genuinely normative norms still play an important role even in variants of Dorsey’s example. Imagine that LaLoosh made a deathbed promise to his father to not get bound up in the norms of baseball excellence. Maybe his father thought he ruined his own life pursuing baseball excellence, and didn’t want his son doing the same. The deathbed promise let LaLoosh play the game, but only as a means to an end. In that case, it’d be *in*appropriate for Davis to blame LaLoosh for not respecting the game. LaLoosh *shouldn’t* respect the rules of the game, because

⁵Now the lesson Darwall himself might draw is that the problem arises from the turn to constitutive norms. Maybe Davis’ blame isn’t *normatively* appropriate – and that’s why baseball norms aren’t normative. I’d be happy to grant that diagnosis. Granting it means conceding that Darwall’s approach can’t incorporate constitutive norms. (I do doubt the diagnosis is right. Davis’ blame looks at least normatively permissible, even if normatively optional, and it’s hard for Darwall’s approach to explain why.)

of his deathbed promise. Maybe Davis could still blame him. But blaming him would be *inappropriate*, once Davis knows all the facts.

Our imagined deathbed promise plays a defensive role: it defends LaLoosh against blame. This paper proposes construing normative facts as those capable of playing that sort of defensive role. The deathbed promise grounds an important fact: that LaLoosh ought not respect the rules of the game. I propose that that fact is normative when it can shield LaLoosh from negative reactive attitudes like blame.

REACTIVE CONSTRUAL (*First Pass*) the fact that x ought to ϕ according to some standard is normative iff it can make negative reactive attitudes about x 's ϕ -ing constitutively inappropriate.

This Construal incorporates a different constitutive norm than the Darwall-Style Construal. It supposes that there is a class of facts with the form ' x ought to ϕ ' that:

- can make blaming x for ϕ -ing constitutively *inappropriate*.
- can make guilt over x 's ϕ -ing constitutively *inappropriate*.
- ...

And it holds that that class of facts just is the class of normative facts. The Darwall-Style Construal focuses on an *offensive* role: metaphorically, normative facts generate swords that we can use to attack others. It centers an offensive role because it takes normative facts to *rationalize* appropriate blame. My Reactive Construal, in contrast, focuses on a *defensive* role normative facts might play: they're *shields* that protect us. It identifies a common core shared with the Darwall-Style Construal because the cases where it recognizes shields are also cases where the Darwall-Style Construal predicts that others can't have swords to attack us. My Reactive Construal thus fits with a richer social picture of negative reactive attitudes that allows them to have a variety of sources – normative and non-normative both.

I'm grouping attitudes like resentment, guilt, and shame under a common label; I called them 'negative reactive attitudes'. P. F. Strawson identified a deep unity that justifies the common label. He takes those attitudes to be:

reactions to the quality of others' wills towards us, as manifested in their behaviour: to their good or ill will or indifference or lack of concern. (Strawson 1974, 15)

You might run over my toe because you didn't care enough to check. I can intelligibly resent what you did, because your action displays indifference to me: my resentment is a reaction to the quality of will directed at *me*. A third party could also be *indignant* that you ran over my toe: a reaction to the quality of will directed at *another*. I'm following Strawson in calling the attitudes that react to displayed quality of will the "reactive attitudes".

So here's another warm-up example to introduce the Reactive Construal. Imagine running over your toe *is* necessary to save a life. Then running over the toe *won't* display negative quality of will. It instead displays moral decency. The fact that I ought to save the life can 'shield' me by preventing my action from displaying a negative quality of will to you. Crucially, though, it's only *some* kinds of facts about what I ought to do that provide that kind of shield. Back to Thomson: if you don't move your bishop horizontally, hundreds will die! The rules of chess say that you ought not move the bishop horizontally. But if you refuse to move the bishop because of the rules of chess, you *are* displaying your quality of will to the hundreds who will die. The rules of chess don't shield you.

2.2 The Official Formulation of the Reactive Construal

I thus follow Strawson and others in taking the reactive attitudes to respond to displayed quality of will. The official version of my Reactive Construal will take the *normativity* of a fact to involve a distinctive capacity to prevent actions from displaying negative quality of will. This version will thus incorporate constitutive norms about negative quality of will.

- **resenting** x's ϕ -ing is constitutively appropriate only if x's ϕ -ing displays a negative quality of will.
- **guilt** about ϕ -ing is constitutively appropriate only if my ϕ -ing displays a negative quality of will.
- ...

And the official formulation of the Reactivity Construal distinguishes normative facts as those capable of 'shielding' actions from displaying negative quality of will.

REACTIVE CONSTRUAL (*Official*) the fact that x ought to ϕ according to some standard is normative iff it can prevent x's ϕ -ing from displaying negative quality of will.

This formulation subsumes the earlier formulations, because negative reactive attitudes constitutively respond to negative displayed quality of will. For example, if a fact prevents an action from displaying an agent's negative quality of will, it'd make resentment constitutively inappropriate.

This proposal avoids circularity by taking the constitutive norms on negative reactive attitudes to *themselves determine* what displays negative quality of will. The moral requirement to save a life at the cost of a toe can make resenting the action that destroyed the toe constitutively inappropriate. It makes resentment inappropriate because of constitutive norms on *resentment*, not constitutive norms on morality. So the Reactive Construal interprets the normativity of morality as an *extrinsic* property of morality: morality is normative *because of* constitutive norms on negative reactive attitudes. This explanatory order avoids circularity: the righthand side of the Reactive Construal doesn't contain anything normative. You might then worry that the Reactive Construal makes the normativity of morality mere sociology. That's a genuine worry that §6 takes up.

The Reactive Construal fits the LaLoosh and Davis example. In the movie, Davis reacts to the quality of will that LaLoosh displays: he doesn't 'respect the game'. However, *had* LaLoosh made a deathbed promise to *not* get invested in respecting the game, his indifference to the game *won't* display his negative quality of will. It'd instead display his moral commitment. That's why the deathbed promise would make blame inappropriate.

The Reactive Construal focuses just on negative reactive attitudes, like resentment and guilt. It's silent about positive reactive attitudes like gratitude. But the focus on negative reactive attitudes is just what we should expect given my initial ambition. My initial ambition is to identify the best contemporary descendant of Mill's account, which focuses on *sanctions*. Gratitude isn't a sanction. Resentment is. So developing a Millian account means focusing on negative reactive attitudes, not positive ones. And the negative ones are the ones we most need shielding against, as Gary Watson explains:

blaming tends to be a much more serious affair; reputation, liberty, and even life can be at stake, and understandably we are more concerned with the conditions of adverse treatment than with those of favorable treatment (Watson 2004, 283)

Given this understandable concern, the constitutive norms on *nega-*

tive reactive attitudes should figure centrally in our social lives, just like the Reactive Construal has it. There may also be interesting connections between normative facts and *positive* reactive attitudes like gratitude. But we wouldn't expect those connections to be as systematic as the connections with the negative ones.

In talking about 'quality of will', I'm following Strawson in labeling a phenomenon that I expect you to already find familiar: what you react to in central cases of resentment or indignation. I'm expecting you to see that running over a toe doesn't display your quality of will if it's necessary for saving a life. I think Strawson is right that "it does not seem to me to matter if a strict definition [of 'reactive' and 'quality of will'] is not to be had" (Strawson 1980, 266); ostension is enough for recognizing the crucial phenomenon.

Another way to describe the Reactive Construal is that it takes genuinely normative facts to constrain our interpretation of behavior. When I think that someone acted as she normatively ought but don't know her motives, I can't appropriately take her action to display negative quality of will. Attempts at mindreading are out of place; guesses about why she 'really' did what she did can't rationalize resentment or indignation. This description of the Reactive Construal makes it easier to see why normativity matters. It helps us navigate our complex social world; we'd be worse off if we stopped taking normative judgments to govern negative reactive attitudes. Then we'd be unpredictably subject to negative emotional sanctions.

3 Plausibility

The Reactive Construal secures the three desiderata on construals of normativity.

3.1 The third desideratum: extensional fit

The Reactive and Practical Construals both explain when a *particular ought* fact is normative. *Domains* are normative derivatively: maybe they're normative when facts in that domain are normative unless trumped by some other domain. In evaluating whether the Reactive Construal is extensionally plausible, we evaluate its verdicts about particular facts. We can't consider all facts, any more than we can when evaluating the Practical Construal.

I break my arm on the way to your party. I ought to go to the ER. You can't appropriately resent my missing your party. My missing it doesn't display negative quality of will to you. It instead

reflects my attitude to myself.⁶ Guilt about missing the party would be similarly inappropriate. The Reactive Construal then classifies this *ought* fact as genuinely normative. This paper's core conjecture about prudence is that many facts about what I prudently ought to do play the role that distinguishes normative facts: shielding me from displaying negative quality of will.

I'm playing a game of chess. If I finish the game, I'll miss the party you invited me to. Given the aim of winning, I ought to make this move here, and that move there. But it'd be appropriate for you to resent my continued play: I do display negative quality of will towards you. Your *appropriate* resentment suggests that those facts about what I ought to do given the aim of winning the chess game aren't normative. If they were, complying with them would shield me from resentment in a trivial case like this. I'm conjecturing that this kind of game-generated *ought* patterns differently than the *oughts* of prudence, with only the latter shielding us. Given that conjecture, the Reactive Construal would agree with the Practical Construal about prudence (normative) and chess (non-normative).

We might then ask *why* the *oughts* of prudence shield us even though the *oughts* of chess don't. Good question! But that's not the question I'm asking in this paper. My only aim here is to identify the normative question – to explain what we're doing when we're asking if prudence or morality is genuinely normative. Even once we've identified the normative question, we still need to answer it. Proponents of the Practical Construal similarly need to answer the normative question, construed Practically. Korsgaard and Gibbard both favor the Practical Construal but answer it differently.

I'm on my way to your party and I see an infant drowning in a puddle. I ought to save the infant, morally speaking. Again, you couldn't appropriately resent my missing the party. My absence doesn't display a negative quality of will to you; it instead reflects my moral decency. The Reactive Construal then classifies this moral *ought* as genuinely normative. It shields me from displaying negative quality of will. Moreover, the actual constitutive norms on resentment seem to prioritize morality over prudence, at least sometimes. Prudence might tell me to not get my shoes messy by saving the baby, while morality disagrees. Our actual constitutive norms treat morality as trumping prudence, at least in this case. The Reactive Construal thus differs from Copp's skepticism about normativity, be-

⁶Maybe I can throw this shield away. Maybe this fact about what's prudent doesn't figure in what motivates me – if I bizarrely go to another party with a broken arm.

cause he holds that there are no normative facts when morality and prudence disagree (Copp 1997).

I hurtfully believe that you're a thief and a liar. You could appropriately resent me for believing so in the absence of evidence. Absent evidence, I'm displaying negative quality of will. But sufficient evidence can mean that I ought to believe it. That epistemic *ought* makes resentment about the belief inappropriate. The Reactive Construal thus promises to classify certain *epistemic oughts* as genuinely normative, as shielding from displaying negative quality of will. Others have defended similar accounts. For instance, Antti Kauppinen links epistemic norms with accountability, suggesting that "the basic way of holding someone epistemically accountable is subtracting credibility points from someone" (Kauppinen 2018, 6).

The Reactive Construal correctly classifies the core cases: it classifies central *oughts* of prudence, morality, and epistemology as all normative. It also classifies a game-playing *ought* as *non*-normative – so too for etiquette. You can appropriately resent me for prioritizing etiquette over you; others can be indignant. Even if local etiquette forbids confronting a casually cruel uncle, you can appropriately resent failures to confront him.

The Reactive Construal also classifies controversial cases as controversial. Philistine that I am, I find aesthetic norms *non*-normative. I'm indignant that Gauguin abandoned his family, and I do not think that his artistic achievements could make my indignation inappropriate. Bernard Williams (1981) finds himself conflicted. If you think that indignation at Gauguin could be inappropriate, you seem to be treating aesthetic *oughts* as normative. Or maybe Gauguin aesthetically ought to break a promise to get lunch with me so that he can work on a masterpiece. If resentment is inappropriate – if he'd be shielded – then the Reactive Construal would classify this aesthetic *ought* as normative. I myself still don't see it: I don't think he would be shielded. If you disagree, and think that he would be shielded, you seem to be treating the aesthetic obligation as normative – just as the Reactive Construal predicts.

The Reactive Construal also allows but doesn't require that legal norms are normative. Some insist that legal norms can fail to be normative (Hart 1982, 267) – and even John Finnis describes the natural law tradition as allowing that deeply unjust legal systems "cannot create what any self-respecting person would count as a genuine obligation" (Finnis 2002, 22-3). He'd think that we can appropriately be indignant at someone who obeyed deeply unjust laws. It's hard to see how someone could disagree with him unless

they saw even deeply unjust legal systems as normatively binding. Shields against negative reactive attitudes figure in the cash value of a genuinely normative obligation, as the Reactive Construal holds.

Now you might worry that etiquette *does* sometimes appear to ‘shield’ us. You might resent me not clapping partway through your performance, as expressing insufficient enthusiasm. You might then learn that the etiquette for classical music forbids applauding during classical music performances. Then you could think that your initial resentment is *inappropriate*. In this case, etiquette does appear to shield me. If so, we’ve found a counterexample to the Reactive Construal: this demand of etiquette isn’t normative, but it does seem to shield me from resentment.

I answer by amending the Reactive Construal to require full information.

REACTIVE CONSTRUAL (*Official*) the fact that x ought to ϕ is normative iff it can prevent x’s ϕ -ing from displaying negative quality of will even given full information.

Resentment based on ignorance is then simply *irrelevant* for the Reactive Construal. I suppress the restriction to full information throughout most of the paper, for simplicity. (I did when presenting the Practical Construal, too.) But restricting attention to full-information cases is legitimate in comparing the Reactive and Practical Construals. Even the most austere versions of the Practical Construal allow that full information matters (Hume 1739; Williams 1979).

This full-information restriction still predicts that etiquette doesn’t shield against resentment. Imagine that everyone around me is breaking the etiquette norms by clapping partway through your performance, but I refuse to join in. At that point, you can appropriately resent me for not joining in even knowing the etiquette for classical music. My refusing to join in shows fetishistic concern for that etiquette. But *if* etiquette were genuinely normative, resentment would be *inappropriate*. Consider morality again. Imagine that everyone else walked by a drowning baby to go to your party. I miss your party because I stop to save the baby. Resenting me for missing the party would still be *inappropriate*.

The pivotal point is that morality *originates* a shield against negative reactive attitudes. Etiquette doesn’t itself *originate* shields against those attitudes, though it can affect how we interpret shields from other sources. I might have a moral obligation to express gratitude. And local norms of etiquette can affect how to express grati-

tude.⁷ The rules of chess similarly determine what's cheating without themselves *originating* shields against displayed will. When I ought to avoid cheating, perhaps because of an implicit promise, following the rules of chess can shield me. But the rules of chess still don't originate a shield against cheating; they only affect how we interpret an independently existing shield.

The Reactive Construal thus allows other considerations to affect the ability of normative facts to shield us. I just mentioned how full information affects that ability. The Construal also allows that *bad* motives can undermine the ability. Someone's delighted to get to hurt me, running over my toe to save a drowning baby; for them, morality's just an excuse. I can appropriately resent what they did. But the Reactive Construal needs only to hold the shield is available given *some* motive for doing the act – not that it's available given *every* motive. That fact is normative because it shields someone who runs over my toe from good motives.

3.2 The first desideratum

I grant that the *Practical* Construal also fits our intuitive sense of what's normative fairly closely. This section briefly explains why we would expect the Practical Construal to fit our intuitive sense even if the Reactive Construal is true.

I care about my interpersonal relationships. That's enough to explain why I care about doing what I normatively ought; failures to do so leave me open to strains on those relationships. For instance, I'm motivated to keep promises because breaking them leaves me open to strains on my relationships – open to my friends and family appropriately resenting or being indignant about the broken promises. I care about their indignation even when the broken promises leave them unaffected. They could appropriately be indignant even if I only break promises to third parties, and I care about avoiding their indignation.

Nor is my concern for interpersonal relationships idiosyncratic. Strawson thinks we all share it:

the human commitment to participation in ordinary interpersonal relationships is, I think, too thoroughgoing and deeply rooted for us to take seriously the thought that a general theoretical conviction [like determinism] might so change our world that, in it, there were no longer

⁷That's how the Reactive Construal would incorporate the moral significance of etiquette that Sarah Buss (1999) nicely describes.

any such things as interpersonal relationships as we normally understand them; and being involved in interpersonal relationships as we normally understand them precisely is being exposed to the range of reactive attitudes and feelings that is in question. (Strawson 1974, 14)

This paper assumes that Strawson is right about the inescapability of these relationships. If he is, my concern for my relationships won't be idiosyncratic. We *all* will care about what we normatively ought to do because we're enmeshed in interpersonal relationships. Pamela Hieronymi distinguishes several different strands of Strawson's account. The most straightforward one is a simple Humean one, where Strawson is claiming that "our psychologies ensure we will not change" (Hieronymi 2020, 47). That simple Humean thought explains why we'd be apt to confuse the Practical and Reactive Construals – and the other strands that Hieronymi identifies would too.

Christine Korsgaard would find my appeal to ordinary interpersonal relationships unsatisfying. She'd insist that it doesn't capture the *intimacy* of the connection between normative judgments and motivation. And indeed, it doesn't capture as intimate of a connection as she wants. A mobster could find himself enmeshed in relationships indifferent to the demands of morality. His friends might even resent him for doing what's moral, taking his moral actions to express a negative quality of will to them. That mobster needn't be motivated to do what he morally ought, given the Strawsonian suggestion I just made. But my ambition is **not** to explain why normative judgments tend to motivate.

My ambition is instead to sketch why we could misconstrue the normative question Practically even if the Reactive Construal is right. The explanation is that there's substantial overlap between the two construals *under ordinary conditions*. Ordinary conditions include the kind of interpersonal relationships that immorality would strain. Given those sorts of interpersonal relationships, we would ordinarily be motivated to do what we morally ought. And we're asking the normative question in our ordinary conditions – a mobster's systematically immoral relationships are extreme and unlike ours. So we'd expect our judgments about the normative question to reflect a link to motivation, since they do in the ordinary conditions where we're forming those judgments. Possibilities where I'm raised by wolves are also not ordinary conditions – we shouldn't see judgments about those possibilities as probative, either.

So I hope that the Reactive Construal undercuts motivational internalism – explaining why normative facts appear linked to moti-

vation, even though the genuinely systematic link runs through the reactive attitudes. But I haven't made good on that hope here. I'm just introducing the possibility by arguing that the Reactive Construal is at least on a par with the Practical Construal.

The Reactive Construal also secures the first desideratum from earlier; it explains why the normative question is substantive. Korsgaard gives a helpful metaphor:

If to have knowledge is to have a map of the world, then to be able to act well is to be able to decide where to go and to follow the map in going there. The ability to act is something like the ability to use the map, and that ability cannot be given by another map. (Nor can it be given by having little normative flags added to the map of nature which mark out certain spots or certain routes as good. You still have to know how to use the map before the little normative flags can be of any use to you.) (Korsgaard 2008, 315)

She's thinking that the ability to use the map is a practical ability – an ability to use the map to *act*. The Reactive Construal instead takes the ability to be a *social* ability – an ability to use the map in navigating our shared social world. Both options explain why little normative flags won't help. The little flags won't help someone who lacks the relevant social ability, on the Reactive Construal, and won't help someone who lacks the relevant practical ability, on the Practical Construal.

Korsgaard discusses a predecessor to the Reactive Construal in discussing Mill. She admits that Mill gives the right sort of answer, but objects that he misses an essential point: “the normative question must be answered in a way that addresses the agent who asks it. And according to Mill's own theory this argument cannot address the agents it is meant for. If they are not utilitarians, it cannot matter to them that utilitarianism would seem normative to people who had been brought up to it” (Korsgaard 1996, 85-6). My Strawsonian spin on the Reactive Construal avoids this objection. Our interpersonal relationships *guarantee* that the Reactive Construal addresses the agents it's meant for, since we're all inescapably susceptible to the reactive attitudes.

3.3 The second desideratum

My core conjecture in this paper is that our intuitive sense of norms that shield us from displaying negative quality of will matches our

intuitive sense of what's normative. The intuitive matches suggest that the Reactive Construal may be an *eligible* candidate for the normative question, like the Practical Construal. The Reactive Construal explains our intuitive sense of what's normative as reflecting constitutive facts about our *actual* reactive attitudes. Maybe, for instance, our convictions about what's normative reflect the fact that Foot's Aristotelianism constitutively governs our actual attitude of resentment, at least in part. That *constitutive* fact about resentment explains the judgments that we make about what's normative.

You might worry that the Reactive Construal is plausibly only for *morality*, not for normativity as such. Shame and guilt may seem like distinctive reactions to moral wrongs in particular. Maybe. (I'm skeptical. I could see a different version of LaLoosh feeling guilty for not respecting the game.) But the Reactive Construal doesn't require the *same* sanctions to always matter. It distinguishes normative facts as shielding an action from reactive sanction – which works so long as resentment, or blame, or guilt, or shame, or ... would otherwise apply. Blame, for instance, seems like enough of a general-purpose sanction for the Reactive Construal. Even if it's not general-purpose enough, *some* kind of reactive sanction will be appropriate for purely self-regarding action, since I can intelligibly sanction myself for actions that are imprudent but morally permissible.

The Reactive Construal explains our intuitive sense of what's normative as reflecting constitutive facts about our *actual* reactive attitudes. David Enoch (2006) might immediately object that Scanlon's contractualism could constitutively govern *schresenment*: actions needn't display *schquality* of will if his contractualism requires them. But that wouldn't be an objection to the Reactive Construal! It's instead evidence that the Reactive Construal secures the second desideratum from earlier. That desideratum was that construals of normativity illuminate when a species has normative thoughts. Maybe we feel resentment and guilt, but encounter another species who feels schresentment and schguilt. That is, the other species also has emotional reactions to behavior that displays negative schquality of will. All that seems like good evidence that they're having normative thoughts! So the Reactive Construal secures the second desideratum from earlier: it illuminates when groups are having normative thoughts.

In saying that the Reactive Construal secures the second desideratum, I'm implicitly extending the Construal. I'm extending it so that it also explains what it is to *treat* a fact as normative. Officially, it's just an account of normativity itself.

REACTIVE CONSTRUAL (*Official*) the fact that x ought to ϕ is normative iff it can prevent x 's ϕ -ing from displaying negative quality of will given full information.

But it also predicts when someone *treats* a fact as normative.

REACTIVE CONSTRUAL (*Treatment*) someone treats the fact that x ought to ϕ as normative iff that person treats it as capable of preventing x 's ϕ -ing from displaying negative quality of will.

Accounts of social phenomenon often predict when someone *treats* the social phenomenon as occurring. For instance, J. L. Austin gives an account of what it is to name a ship – and that account also predicts when someone competent would recognize a ship-naming. Since the Reactive Construal treats normativity as a social phenomenon, it should also predict when someone *treats* a fact as normative.

4 Defending parity

This paper aims to establish parity between my Reactive Construal of normativity and the Practical Construal – establishing that the Reactive Construal deserves serious consideration too. Establishing parity doesn't mean showing that the Reactive Construal is *true*, since many doubt that the Practical Construal is true.

Some objections to the Reactive Construal rest on misunderstandings. The Construal links normative facts with appropriate reactive attitudes. False normative judgments are irrelevant, no matter how widely accepted. For instance, Kate Manne emphasizes the significant moral mistakes encoded in our dispositional reactive attitudes, like dispositions to resent women for not providing gendered goods (Manne 2017, 147). We should see resentment, guilt, or shame about such refusal as inappropriate. In evaluating the Reactive Construal, focus on the normative facts in asking when you'd find the reactive attitudes appropriate.

Now the Reactive Construal does force distinctive conclusions about dilemmas. I unwittingly promised to meet two people for lunch today; I can't meet both. The Construal predicts that there's *no particular action* I normatively ought to do in dilemmas, since someone's resentment remains appropriate whatever I do. The Practical Construal makes the same prediction: there's no particular action I ought to do, since there's no one action that a rational person would be motivated to do. However, the Practical Construal

does naturally capture *graded* comparisons: missing both lunches is normatively worse than missing just one, because a rational person would be motivated to go to at least one lunch. The Reactive Construal struggles a bit to make graded comparisons. Maybe missing both lunches is normatively worse because *more* instances of resentment would be appropriate: resentment for missing the first lunch as well as missing the second. The graded comparisons ground the normative fact that I ought to go to at least one lunch. I suspect that the Reactive and the Practical Construals should both start by characterizing when one action is normatively worse than another. But I couldn't write this paper about graded normativity, because extant proponents of the Practical Construal tend to not see graded normativity as central.

Nicholas Southwood objects that views like the Reactive Construal are objectionably parochial. He cites “evidence that the tendency to experience guilt, blame, and resentment and to regard them as fitting is culturally specific (Benedict 1946, Dodds 1951). [Such views] would therefore seem to have the unpalatable consequence that the members of certain cultures are incapable of making genuine moral [or normative] judgements” (Southwood 2011, 769). This objection is pressing. Since the Reactive Construal aims to capture normativity, it would be refuted if some cultures make normative judgments that it can't capture.

Strawson himself includes “the more complicated phenomenon of shame” (Strawson 1974, 16) among the reactive attitudes. He includes it because shame can be a reaction to negative quality of will; I can feel shame because I'm the sort of person who displays disregard for you. However, I can also feel shame for other reasons, about bodily functions, say – and that doesn't have anything to do with negative quality of will. I'll call the former kind of shame ‘reactive shame’. The Reactive Construal takes reactive shame to determine what members of a shame culture are treating as normative. The fact that I ought to ϕ would be normative when it makes *reactive shame* constitutively inappropriate.

All accounts of shame allow for reactive shame. Some accounts of shame center its audience; for instance, Stephen Darwall contrasts what he understands as *second-person* responses like guilt and *third-person* responses like shame. For him, shame involves “seeing oneself as an object of the other's regard or “gaze” –of her disdain, perhaps, or of her just seeing through one's public persona to something one is ashamed to have seen” (Darwall 2009, 71). Guilt and resentment, in contrast, involve second-person address. Shame can be a third-

personal reaction to negative quality of will. After all, I also care about how other people talk about me when I'm not in the room – which can involve caring about whether they think I've done something shameful.

Shame is a more holistic response than guilt, as Heidi Maibom explains: “by contrast to guilt, which focuses on the action performed, shame focuses on the self as a whole. Even when shame is precipitated by an action, it is not the action, as such, but being the sort of person who would perform such an action, that the person is ashamed of” (Maibom 2010, 568). Reactive shame makes perfect sense even given this point. In feeling reactive shame, I'm feeling shame about being the sort of person who displays negative quality of will. For instance, I might feel shame about being the sort of negligent person who'd run over your toe because I wasn't paying attention. And that sort of shame involves the reaction to displayed will that distinguishes reactive shame.

The Reactive Construal takes reactive shame to distinguish normative facts in shame cultures: normative facts constitutively shield against *that kind* of shame. So the Construal isn't a parochial artifact of guilt cultures. It gives the exact same account of normativity in shame cultures. It focuses on a class of emotions that play a common functional role: the role of reacting to negative quality of will. Reactive shame illustrates how different reactive sanctions can play that functional role.

You may find that your grasp of reactive sanctions depends on your grasp of what normatively ought to be done. For instance, you might only understand distinctively reactive shame as the special kind of shame that requires the judgment that the agent normatively ought not to do it. Or I might blame a chair for being in my way if I stub my toe. But that's not the distinctively *reactive* kind of blame that Strawson and I are discussing. If your grasp on reactive sanctions runs through normative judgments, then the Reactive Construal is true. The Construal ties normative oughts with constitutive facts about reactive attitudes like resentment and shame.

But isn't the link between reactive sanctions and normativity then circular? Our understanding of what's normative then tells us which sanctions are reactive. I answer that our grasp of a phenomenon needn't mirror its metaphysical structure. In this case, our grasp of reactive sanctions (like reactive shame, or reactive blame) might depend on our grasp of what's normative, even though constitutive facts about reactive sanctions determine what's normative. That is, constitutive facts about reactive sanctions *appear to us as*

facts about what's normative – we fail to realize that those constitutive facts are explaining our grasp of what's normative. Here's a companion in innocence. Tarski (1933) gave a model-theoretic account of logical consequence and logical truth. But our grasp on the model theory depends in part on our grasp of logical consequence, since the model theory assumes a grasp of which vocabulary is logical and which isn't. But the metaphysical structure of Tarski's account is independent from our grasp of logical consequence, since that account classifies certain set-theoretic truths as logical truths even though we don't intuitively see them as logical truths (Etchemendy 1990).

5 Isn't the Reactive Construal mere sociology?

I close with a natural worry: that the Reactive Construal makes normativity mere sociology. It appears to reduce normative authority to how we happen to treat each other. But normative authority is more: it's a standard for how we *should* treat each other that's separate from what we actually do. Even worse, the Reactive Construal suggests the wrong methodology for determining what's normative. It suggests that we look to sociology!

Let's back up. We're trying to understand the question 'why be moral?' My grip on this question runs through my grip on paradigmatic examples. Prudence contrasts with the rules of chess; I should act prudently, normatively speaking. In contrast, it's not true that I should obey the rules of chess, normatively speaking. When we're asking 'why be moral', we're looking for similarities between prudence and morality, or between the rules of chess and morality.

The Reactive Construal could identify the real similarity between prudence and morality. They're both normative because they govern reactive sanctions. This identification isn't sociology at all. It rather identifies the constitutive norms that help make negative reactive attitudes what they are. And sociological inquiry into how we happen to treat each other cannot reveal constitutive norms. Instead, we learn about constitutive norms 'from the inside' – from our own reflection on the practices we inhabit. The Reactive Construal *retains* a standard methodology for determining what's normative. It starts with examples where I'm most confident – prudence, say – and tries to characterize those examples. Then it moves to more difficult examples, like morality or the law. The objects of inquiry are the facts about prudence and the facts about morality.

The Reactive Construal helps some philosophers who reject the Practical Construal. Philippa Foot is one example – she questions

“whether it is right to think that moral action has to be brought under a pre-established concept of practical rationality” (Foot 2001, 10) that links it to motivation. She suggests that we flip our ambitions: start by taking morality to constrain what’s rational, rather than starting with an independent conception of what’s practically rational. Her suggestion is the covert motivation for this paper. I’m setting out an alternative picture of practical rationality that doesn’t incorporate the tight link to motivation that the Practical Construal does.

Foot herself ends up dismissing the question ‘why be moral?’:

If the sceptic ... still goes on saying that he has not been shown that there is reason for acting as a good person would act, it is no longer clear what he is asking for. To ask for a reason for acting rationally is to ask for a reason where reasons must a priori have come to an end. And if he goes on saying ‘But why should I?’, we may query the meaning of this ‘should’ (Foot 2001, 65)

But these ‘should’s arise spontaneously – the undergraduates generate them after reading ‘Famine, Affluence, and Morality’, and Glaucon did too. I find this passage uncomfortably reminiscent of passages from Ayer, where he dismisses questions about God or justice. The queries target central questions that pull people into philosophy.

The Reactive Construal lets Foot formulate the meaning of this ‘should’ without abandoning her guiding insight. On the Reactive Construal, the skeptic is asking why moral obligations can shield us from resentment and indignation in the way that prudential obligations can. An answer would explain why acting from moral obligations and acting from prudential obligations *does* shield us from displaying a negative quality of will. Foot herself can explain the pivotal constitutive norms on reactive sanctions as grounded in the life-form of our species. This extension is perfectly consistent with Foot’s other commitments, while rendering the skeptic’s question intelligible. The Reactive Construal frees her to identify a deeper real similarity between morality and prudence: constitutive norms on reactive sanctions arising from the life-form of our species.

This paper hasn’t argued that the Reactive Construal is better than the Practical Construal. It has only aims at establishing parity between the Reactive Construal and the Practical Construal. The Practical Construal structures vast swathes of metaethics – enough so that Michael Smith feels entitled to call a question it generates “the” moral question (1994)! The Reactive Construal illustrates how

metaethics could be different. If it's right, we should replace tired discussion of realism and motivation with better discussion of realism and displayed quality of will. And even if it's wrong, it advances our understanding of central questions in metaethics, because it helps us identify the fixed points on adequate construals of normativity. It provides a concrete alternative to the Practical Construal. And if we learn that that alternative is wrong, we're learning about fixed points of adequate construals of normativity that it doesn't secure.

References

- Baker, Decker. 2018. "Skepticism about Ought Simpliciter." In *Oxford Studies in Metaethics* 13, 230–252.
- Blackburn, Simon. 1998. *Ruling Passions*. Oxford: Oxford University Press.
- Buss, Sarah. 1999. "Appearing Respectful: The Moral Significance of Manners." *Ethics* 109 (4):795–826.
- Copp, David. 1997. "The Ring of Gyges: Overridingness and The Unity of Reason." *Social Philosophy and Policy* 14 (1):86–101.
- . 2000. "Milk, Honey, and the Good Life on Moral Twin Earth." *Synthese* 124:113–137.
- . 2020. "Normative pluralism and skepticism about ought simpliciter'." In Ruth Chang and Kurt Sylvan (eds.), *The Routledge Handbook of Practical Reason*, 416–437. Routledge.
- Darwall, Stephen. 2009. *The Second-Person Standpoint*. Harvard University Press.
- . 2017. "What Are Moral Reasons?"
- Dorsey, Dale. 2016. *The Limits of Moral Authority*. Oxford: Oxford University Press.
- . 2020. "Respecting the Game: Blame and Practice Failure." *Philosophy and Phenomenological Research* 101 (3):683–703.
- Eklund, Matti. 2020. "The Normative Pluriverse." *Journal of Ethics and Social Philosophy* 18 (2):121–146.
- Enoch, David. 2006. "Agency, Shmagency: Why Normativity Won't Come from What Is Constitutive of Action." *The Philosophical Review* 115 (2):169–198.
- . 2011. *Taking Morality Seriously: A Defense of Robust Realism*. Oxford: Oxford University Press.
- Etchemendy, John. 1990. *The Concept of Logical Consequence*. Cambridge, MA: Harvard University Press.
- Finlay, Stephen. 2019. "Defining Normativity." In *Dimensions of Normativity: New Essays on Metaethics and Jurisprudence*, 187–219. Oxford University Press: Oxford.

- Finnis, John. 2002. "Natural Law: The Classical Tradition." In *Handbook of Jurisprudence and Philosophy of Law*, 1–60. Oxford University Press: Oxford.
- Foot, Philippa. 1972. "Morality as a System of Hypothetical Imperatives." *Philosophical Review* 81(3):305–316.
- . 2001. *Natural Goodness*. Oxford: Oxford University Press.
- Gibbard, Allan. 1990. *Wise Choices, Apt Feelings*. Cambridge, MA: Harvard University Press.
- Hart, H. L. A. 1982. *Essays on Bentham*. Oxford University Press: Oxford.
- Hieronymi, Pamela. 2020. *Freedom, Resentment, and the Metaphysics of Morals*. Princeton, NJ: Princeton University Press.
- Horgan, Terence and Timmons, Mark. 1996. "From Moral Realism to Moral Relativism in One Easy Step." *Critica* 28:3–39.
- Hume, David. 1739. *A Treatise of Human Nature*. Oxford: Clarendon, 2nd edition.
- Kauppinen, Antti. 2018. "Epistemic Norms and Epistemic Accountability." *Philosophers' Imprint* 18 (8):1–16.
- Korsgaard, Christine. 1996. *The Sources of Normativity*. Cambridge: Cambridge University Press.
- . 2008. "Realism and Constructivism in Twentieth Century Moral Philosophy." In *The Constitution of Agency: Essays on Practical Reason and Moral Psychology*. Oxford University Press: Oxford.
- Maibom, Heidi. 2010. "The Descent of Shame." *Philosophy and Phenomenological Research* 80 (3):566–595.
- Manne, Kate. 2017. *Down Girl: The Logic of Misogyny*. Oxford University Press: Oxford.
- Mill, John Stuart. 1863. "Utilitarianism." In Russ Shafer-Landau (ed.), *Ethical Theory: An Anthology*. Malden, MA: Wiley-Blackwell;.
- Parfit, Derek. 2011. *On What Matters*, volume 2. Oxford: Oxford University Press.

- Prichard, H. A. 1912. “Does Moral Philosophy Rest on a Mistake?” *Mind* 21 (81):21–37.
- Sagdahl, Mathias. 2014. “The Argument from Nominal/Notable Comparisons, Ought All Things Considered,’ and Normative Pluralism.” *The Journal of Ethics* 18(4):405–25.
- Scanlon, T M. 1998. *What We Owe to Each Other*. Cambridge, MA: Harvard University Press.
- Setiya, Kieran. 2021. “What is morality?” *Philosophical Studies* .
- Singer, Peter. 1972. “Famine, Affluence, and Morality.” *Philosophy and Public Affairs* 1 (3):229–243.
- Smith, Michael. 1994. *The Moral Problem*. Blackwell.
- Southwood, Nicholas. 2011. “The Moral/Conventional Distinction.” *Mind* 120 (479):761–802.
- Strawson, P F. 1974. “Freedom and Resentment.” In *Freedom and Resentment and other essays*. Methuen and Co, Ltd.
- . 1980. “Replies.” In *Philosophical Subjects: Essays Presented to P. F. Strawson*, 260–296. Oxford: Clarendon Press.
- Svavarsdottir, Sigrun. 1999. “Moral Cognitivism and Motivation.” *The Philosophical Review* 108:161–219. ISSN 00318108. doi: 10.2307/2998300.
- Tarski, Alfred. 1933. “The concept of truth in the languages of the deductive sciences.” In *Prace Towarzystwa Naukowego Warszawskiego, Wydział III Nauk Matematyczno-Fizycznych* 34.
- Thomson, Judith Jarvis. 2008. *Normativity*. Open Court.
- Tiffany, Evan. 2007. “Deflationary Normative Pluralism.” *Canadian Journal of Philosophy* 37 (5):231–262.
- Watson, Gary. 2004. “Two Faces of Responsibility.” In *Agency and Answerability*, 260–288. Oxford University Press: Oxford.
- Williams, Bernard. 1979. “Internal and External Reasons.” In Ross Harrison (ed.), *Rational Action.*, 101–113. Cambridge University Press.
- . 1981. *Moral Luck*. Cambridge: Cambridge University Press.