

An epistemology for moral naturalists*

(9,873 words)

Natural sciences like physics, chemistry, and biology are impressively successful. Their success convinces many of us that there are physical, chemical, and biological facts that are independent of and more fundamental than our methods for discovering them. Moreover, it's natural to take their success to suggest a general strategy for establishing a domain D of facts as independent of and more fundamental than our methods for investigating them: show that the D-facts are saliently similar to physical, chemical and biological facts. Since we say that physics, chemistry, and biology are *natural* sciences, it's natural to classify proponents of this strategy as *naturalists* about the D-facts.

Many of us believe that *moral* facts are independent of and more fundamental than our methods for discovering them. The naturalist strategy is an attractive way to vindicate this belief. It's thus unsurprising that recent metaethics has seen a broad range of naturalist views. One way to be a naturalist is to hold that the *real definition* of moral properties includes only natural properties. Another way to be a naturalist is to hold that moral properties are *reducible* to natural properties. And yet another way is to hold that moral properties are *composed of* natural properties. And there are other options, too.¹ There are also lots of options about *which*

*For helpful comments and conversations, I'm grateful to Mike Ashfield, Rima Basu, Selim Berker, Stephen Bero, Renee Bolinger, Sam Carter, Ruth Chang, Heather Demarest, Alex Dietz, Jonathan Drake, Eric Encarnacion, Maegan Fairchild, August Gorman, Michael Hannon, Joe Horton, Nathan Howard, Dan Immerman, Robin Jeshion, Dan Korman, Tanya Kostochka, Nicholas Laskowski, Woo Ram Lee, Ben Lennertz, Alida Liberman, Michael Milona, Abelard Podgorski, Jon Quong, Robert Rupert, Kenneth Silver, Keshav Singh, David Sobel, David Sosa, Sungwoo Um, Aness Webster, Preston Werner, and Jonathan Wright. I'm especially grateful to my commentators at the 2020 Central, Joshua Schechter and Sam Filby, to Julia Staffel and Brian Talbot for very helpful comments on the penultimate version of this paper, and to John Hawthorne, Steve Finlay, Barry Schein, Scott Soames, Ralph Wedgwood, and Mark Schroeder.

¹I intend to be ecumenical in my characterization of naturalists, including both *reductive* and *non-reductive* naturalists; for illuminating discussion, see Dowell (2013) and McPherson (2015).

natural facts matter: facts about our *desires*,² or facts about *welfare*,³ or teleological facts about the *natural life form of our species*.⁴ This paragraph describes moral naturalism as centrally a thesis about moral *metaphysics*. It's an account of what moral facts are, and of why they are independent of and more fundamental than our methods for discovering them.

Unfortunately, though, it's hard for moral naturalists to explain our ordinary attributions of moral knowledge, for a range of reasons we'll encounter as the paper goes on. But our ordinary attributions of moral knowledge seem generally correct. If naturalists can't explain them, we seem to have good evidence against moral naturalism. This paper shows how a new view – what I'll here call “deflationism about moral knowledge” – allows moral naturalists to explain away that apparent evidence. Deflationism has been defended at length elsewhere (Perl 2020b, ms; Perl and Schroeder 2019). My goal here is to show that, if deflationism is true, moral naturalists can cleanly explain our ordinary attributions of moral knowledge.

1 Apriority

In ordinary contexts, we allow that much moral knowledge is available just by thinking, without carrying out empirical investigation. For instance, we take knowledge that my promising to ϕ is a reason to ϕ to be available in that way; it's knowable *apriori*.

Moral naturalists struggle to explain the range of moral knowledge that's plausibly *apriori*. A range of philosophers worry that moral naturalism is false because it can't explain the *apriority* of moral knowledge (Huemer 2005; Laskowski and Howard 2019; Parfit 2011; Rawls 1975; Scanlon 2014; Shafer-Landau 2003, 2006; Wedgwood 2007). In fact, naturalists struggle even to explain how we can know *apriori* that moral rightness isn't a yellow rose (Laskowski and Howard 2019, 8). A recent survey notes that:

the popularity of this kind of naturalism in metaethics seems to have faded. ... A primary problem is the already-mentioned and generally acknowledged intuitionist or apriori epistemology of a robust range of normative truths, which is a primary motivation for Moore and other primitivists' denials that normative properties could be 'naturalistic' (Laskowski

²As Mark Schroeder (2007) suggests.

³As Peter Railton (1986) and Richard Boyd (1988) suggest.

⁴As Philippa Foot (2001) and Michael Thompson (2012) suggest.

and Finlay 2017, 6)⁵

This section briefly explains why *aprioricity* looks like a problem for moral naturalists. As you go through it, you may notice strategies to disarm the problem. And indeed, the rest of the paper will develop my own deflationist strategy for disarming the problem. But I'll claim that my strategy is a *uniquely principled* way to disarm the problem. I've argued in other work that all moral realists should accept deflationism. If I'm right, it won't matter if there are also other strategies for disarming the problem.

It's hard for moral naturalists to explain the *aprioricity* of moral knowledge because naturalists seem forced to explain *apriori* moral knowledge as knowledge of something analytic. If our knowledge isn't knowledge of something analytic, it'd have to be knowledge directly of the natural facts, which seems only available *aposteriori*. And, unfortunately, the current options for explaining our *apriori* knowledge as knowledge of something analytic face insuperable problems.

Options for explaining our *apriori* knowledge as knowledge of something analytic differ about what's analytic. Some, like David Lewis, only take very basic truths to be analytic – that “something of the appropriate category is a value if and only if we would be disposed, under ideal conditions, to value it” (Lewis 1989, 113), where he cashes out the ideal conditions in a naturalistically acceptable way. Truths about promises don't end up analytic, but rather follow from analytic truths together with further empirical truths. But since truths about promises follow only given further empirical truths, the truths about promises aren't themselves knowable *apriori*. So a view like Lewis' ends up conflicting with our conviction that we can know about the moral significance of promises *apriori*.

So it seems like a naturalist who wants to vindicate our ordinary convictions must take a much wider range of truths to be analytic. And indeed that's just what Frank Jackson (1998) suggests. Our *apriori* knowledge about promises would be knowledge of analytic truths about 'reason' (or perhaps about 'promise'.) A core problem for this idea is that *utilitarians* also seem to be semantically competent with those words and yet deny that facts about promises are (objective moral) reasons. So if utilitarians and defenders of common sense like Ross are both semantically competent, they must be competent with subtly different things: 'reason_U' and 'reason_R'. Since the words are different, both parties could be correct: 'promises aren't reasons_U' could be analytically true even while 'promises are reasons_R' is also analytically true.

⁵Compare too Pigden (2012) .

And the conclusion that utilitarians and Rossians are both right seems too high a price. For one thing, that conclusion seems to break the analogy with natural science that opened this paper. We wouldn't want to say that Newtonian mechanics and relativistic mechanics are both right; saying so seems to give up our conviction that they're both talking about the same thing. Relatedly, a Jackson-style view seems to predict that (*) is true:

(*) Rossians know that promise-breaking is sometimes wrong just because it's promise-breaking, while utilitarians know that promise-breaking is never wrong just because it's promise-breaking.

Rossians are thinking about 'wrong_R', and utilitarians about 'wrong_U'. Since they're thinking about different things, they can both know what they believe. Now Jackson himself retreats to suggestions about what would be analytic in 'mature' theory. And that retreat seems to abandon his naturalist ambitions, as Yablo (2000) explains.

Analytic naturalism thus seems to face an insuperable dilemma. Either truths about promises are analytic or they're not. If they're not, analytic naturalists can't explain the full range of *apriori* moral knowledge. But if truths about promises *are* analytic, morality is very much unlike natural science: apparently inconsistent views like utilitarianism and Rossianism can both be true because they're talking about subtly different subjects.

This paper provides a new account that allows moral naturalists to do better: to explain the *aprioricity* of moral knowledge without objectionable costs. Moreover, the new account has several other advantages in moral epistemology, dissolving other important challenges for moral naturalists. Now I do acknowledge that some naturalists are serenely indifferent to explaining our ordinary attributions of moral knowledge. Maybe they think that empirical methods are the only methods that deliver genuine knowledge. So *part of what they care about* in developing their naturalist view is showing how moral knowledge can be grounded in empirical methods. I'm not talking about those naturalists.

I'm instead talking to those who find moral naturalism attractive centrally for its metaphysical upshots: those who find it attractive as a way of vindicating our conviction that moral facts are independent of and more fundamental than our methods for investigating them, and as a way of explaining:

- why moral properties supervene on natural properties (cf. Brown

(2011); Dreier (1992); Dunaway (2015); Jackson (1998); Ridge (2007); Schroeder (2005));

- how moral properties fit in the natural world. (cf. Boyd (1988); Brandt (1979); Brink (1989); Copp (2007); Foot (2001); Hursthouse (1999); Jackson (1998); Railton (1986); Schroeder (2007); Sturgeon (2006))
- how moral terms manage to refer to moral properties. (cf. Boyd (1988); Brink (2001); Copp (2000); Dowell (2016); Dunaway and McPherson (2016); Jackson (1998); Schroeder (ms)).

To put my cards on the table: I'd be disappointed if moral naturalism has distinctive upshots for moral epistemology. Those consequences distract from the central virtues, which are in moral metaphysics.

2 The ingredients of deflationism

This section introduces my account of what it takes to have moral knowledge, which I'll call *deflationism* about moral epistemology. Think of deflationism as a novel hypothesis about our ordinary attributions of moral knowledge – that those attributions track an interesting and unappreciated quirk of our concept of knowledge, rather than revealing anything about the nature of morality. I don't expect to convince you here that the deflationist hypothesis is true. Rather, I'll focus on showing that, *if* the deflationist hypothesis is true, *every* moral naturalist can explain our ordinary attributions of moral knowledge. For instance, every naturalist can explain why we take moral knowledge to be available *apriori*.

I'll briefly indicate why we should take the deflationist hypothesis seriously. I myself think that there is *decisive* evidence that it's true. Unfortunately, though, appreciating the decisive evidence takes significant work. I'll focus here just on showing that the deflationist hypothesis about knowledge is *philosophically* very significant, in hopes of motivating you to work out for yourself whether you think it's actually true. This paper is thus in part an extended argument for slogging through my (2020b) or (ms).

2.1 Background: Kratzer

The evidence for deflationism is in the first place linguistic: that it's the best way for moral realists to accept the contemporary orthodoxy that Angelika Kratzer (1977, 1981, 2012) has pioneered. In introducing

the orthodox account, I'll focus on all-things-considered moral judgments, which we express with modal verbs like *may*, *must*, and *have to* and modal adjectives like *wrong* and *permissible*.⁶ Kratzer emphasizes that we use those modals for a range of topics: the demands of morality as well as the demands of prudence and the demands of positive law. Consider (1).

(1) You have to give Bill a hundred dollars.

(1) can be used to express a *moral* obligation to give Bill the money. (Maybe I promised the money to him.) Or it can be used to express a *prudential* obligation. (Maybe giving him the money is the only way to achieve one of my goals.) Or it can be used to express a *legal* obligation. (Maybe a judge ruled that I owe him the money.) Other modal terms are similar. They can be used to express many different kinds of deontic facts. Moreover, modal terms behave remarkably similarly across different natural languages. If a modal term can be used to express one kind of obligation or permission, it can usually be used to express other kinds of obligations or permissions as well.⁷ We should expect the best semantics for modal terms to explain the range of uses.

Kratzer provides the standard unifying account. She suggests that modals like *have to* have at least one more argument-place than their surface syntax suggests. The argument-place is saturated with something that represents the demands of morality, or the demands of prudence, or whatever other kind of demands are contextually salient. It's standard to call her approach a *contextualist* semantics for modal terms, since she takes the context to supply arguments that combine with the semantics for the modal term to determine a complete, truth-evaluable proposition. Now the details of Kratzer's view are very complicated, in ways that are irrelevant at this point. We can abstract away from the complications by saying that the context supplies an *ordering* of worlds that ranks possible worlds in relation to each other.⁸

An ordering-centered semantics like Kratzer's deserves to be uncontroversial because it elegantly explains why *may* and *must* are duals. They're duals in the sense that $\lceil \text{Must}(p) \rceil$ is true iff $\lceil \neg \text{May}(\neg p) \rceil$ is true. For instance, we'll infer (2) from (3), and (3) from (2).

⁶My official account will also generalize to 'reason' and 'good', though I won't generalize it here.

⁷For some exceptions, see Viebahn and Vetter (2016).

⁸Her account actually appeals to ordering sources, which are even more complicated.

- (2) You may not let the drowning child drown.
- (3) You must save the drowning child.

We're ready to make these inferences because we see that (2) is true iff (3) is true.

Kratzer can crisply explain why *may* and *must* are duals. For her, $\lceil \neg \text{May}(\neg p) \rceil$ is true iff it is not true that $\neg p$ is true at some highest-ranked world. But if $\neg p$ is not true at any highest ranked world, then p is true at every highest ranked world. And $\lceil \text{Must}(p) \rceil$ is true iff p is true in every highest-ranked world. So for her, the modals are duals because the quantifiers " \exists " and " \forall " are duals. So she *explains* the duality of modals, rather than just stipulating it.

Kratzer's evidence for ordering-centered *semantics* is also evidence for an ordering-centered account of our *concepts*. Focus on the thoughts that (2) and (3) are used to express, rather than the sentences themselves. We'll think that someone who accepts one thought is also rationally committed to accepting the other one. And they're rationally committed to accepting the other one because the first thought is true iff the second thought is. We can explain the shared truth-conditions by taking the concept CAN and the concept MUST to include an argument-place for an ordering.

I'll here assume that modal terms and modal concepts do have an argument-place for something like an ordering. This assumption is an empirical claim about our talk and thought, not a question mere introspection can settle. But it is highly plausible. In fact, it's the most plausible part of Kratzer's approach, accepted even by those who reject many other parts of the approach. For example, it's incorporated into dynamic approaches (say, by Malte Willer (2014) and also relativist approaches (say, by Niko Kolodny and John MacFarlane (2010))).

2.2 Knowledge about orderings

This paper addresses an intramural dispute among moral realists, about whether fundamental moral properties are natural properties. All moral realists should agree that, given a Kratzerian semantics, there will be some ordering that is the objective moral standard. If some kind of act consequentialism is true, an act consequentialist ordering is the moral standard: that ordering classifies an action as required iff it produces the best outcome, permitted if it produces one of several best outcomes, and otherwise forbidden.

The property *being-the-moral-standard* is necessary for formulating important metaethical debates. Without it, Kratzer's semantics

would take the fact that I have to give Bill a hundred bucks to just be the fact that I give Bill a hundred bucks in all the worlds that some ordering ranks highest. And orderings are just rankings of possible worlds with respect to each other, or sets of propositions that generate rankings of possible worlds, or something similar. It's then hard to see what moral naturalists and non-naturalists are arguing about. They're not disagreeing about the nature of rankings on possible worlds. Nor are they disagreeing about the nature of sets of propositions.

Fortunately, though, we can formulate the debate about moral naturalism by extending Kratzer's approach. The parties to the metaethical debate are not disagreeing about the nature of a Kratzerian ordering. Nor are they necessarily disagreeing about which Kratzerian ordering is the moral standard; they might agree on every normative question. (Think of Sidgwick and Railton, roughly.) They are instead disagreeing about the *explanation why* that Kratzerian ordering captures the demands of morality. The naturalist thinks the explanation is that the ordering has the *natural* property/ properties that constitute being the demands of morality, and the non-naturalist thinks that the ordering has the *non-natural* property/ properties that constitute being the demands of morality.

So naturalists and non-naturalists can both accept Kratzer's semantics by positing further normative facts, and taking a particular Kratzerian ordering to capture the demands of morality because it has the right kind of relationship to those normative facts. I use the predicate "is the moral standard" to refer to the property of having the right kind of relationship to those further normative facts. So I will say that naturalists and non-naturalists disagree about the nature of the property *being-the-moral-standard*. The naturalist takes the property to be a fully natural property, and the non-naturalist disagrees. I earlier assumed that you understood this property: I talked about the ordering that was the objective moral standard, which I take to appeal to this property.⁹

A crucial part of my deflationist idea will be that someone could know that killing is wrong according to a particular ordering. In fact,

⁹Some readers might remain unconvinced that this property figures in what's communicated. If it doesn't, we end up with a dramatically simplified version of deflationism. Such readers should take the §2 account of knowledge of orderings to *fully* ground *apriori* moral knowledge for anyone who has true moral beliefs. Further questions about whether people with such knowledge are lucky in a way that excludes knowledge *can't even be formulated* without talking about the property *being-the-moral-standard* – the only kind of luck that might exclude knowledge is luck in matching the moral standard.

someone could know that killing is wrong according to an ordering that turns out to be the moral standard. Someone could have that knowledge if they have dispositions to reason in ways that match it. For instance, I could know of the ordering x that is the moral standard that killing is usually wrong $_x$ if I had the relevant dispositions about killing: if I'll see killing as usually wrong $_x$, but permissible $_x$ if necessary for self-defense, but not given a non-lethal alternative, and so on. This kind of knowledge is just knowledge about what's wrong according to that ordering. That knowledge is crucially different from further knowledge that that ordering is the objective moral standard. Dispositions to reason correctly about what's wrong according to an ordering don't themselves support *that* further knowledge.

The idea that the right dispositions can support knowledge about the structure of the ordering should already look familiar. A similar idea seems natural about *logic* as well: the right dispositions about logic support knowledge about logical constants, like classical disjunction. If I'm stably disposed to accept that *Disjunction Introduction* is true (that is, that p entails p or q) without error, I'm in a position to know that p or q follows from p .

And I can know some things about an ordering even if I don't know others, just like I can know *some* logical facts without knowing others. For example, I can know that *Disjunction Introduction* is true even if I don't know that *Or-to-If* is true (that is, that not- p or q entails if p , q). I can know that *Disjunction Introduction* is true if I have the relevant inferential dispositions, and I can have those inferential dispositions even if I don't have the inferential dispositions that constitute recognizing *Or-to-If*. More generally, my dispositions can give me *some* knowledge about disjunction even if they don't give me perfect knowledge. Similarly, my dispositions about the moral wrong of killing can support knowledge about an ordering even given ignorance of other parts of the ordering.

Deflationism will hold that propositions about the ordering that's the moral standard will be analytic and so knowable *a priori* for the person with the right dispositions. It thus resembles Jackson's moral functionalism. But it has additional features that distinguish it from Jackson's moral functionalism, and I'll spend the next ten pages describing those additional features. So I won't talk again about *aprioricity* again for about ten pages. But the intervening pages serve to introduce the deflationist account that *preserves* this account of the *aprioricity* of moral knowledge while avoiding important problems.

You should already worry that people with the right dispositions are in an important sense *lucky*, in a way that prevents them from

having genuine knowledge. In developing your worry, you might distinguish at least two kinds of luck. To introduce the first kind of luck, think of someone who is unusually good at math: they can reason out the answers to questions that would stump most of us. They're lucky to have their mathematical abilities. Duncan Pritchard calls this kind of luck 'capacity epistemic luck' (Pritchard 2005, 134), because it's luck that the agent is capable of having knowledge. That luck doesn't seem to undermine knowledge; the person who's unusually good at math does in fact have knowledge. Similarly, you might think that someone with the right inferential dispositions about an ordering has *capacity* epistemic luck: they're lucky to possess the ability that they have, and that luck doesn't undermine their knowledge about that ordering.

However, the person who has the right dispositions about the ordering that's the correct moral standard has another kind of luck. Contrast that person with another person who has dispositions that are a perfect guide to another ordering. We might credit both individuals with knowledge about the relevant orderings. And we might note that both of them have *capacity* epistemic luck, in having dispositions that allow them to know about their ordering. But we might also say that the first one has another kind of luck: luckily enough, the ordering that they know about is in fact the moral standard! That further luck looks different from capacity luck, and it looks much more like the kind of luck that undermines knowledge. So we can reasonably ask whether that further kind of luck undermines knowledge.

3 Deflationism

My deflationist framework predicts that knowledge about orderings can ground genuine moral knowledge. Deflationism holds, roughly, that that knowledge of an ordering can ground moral knowledge for someone who *presupposes* that that ordering is the moral standard. It'll thus vindicate our *apriori* moral knowledge by claiming, very roughly, that *apriori* knowledge of an ordering supports *apriori* moral knowledge for someone who presupposes that that ordering is the moral standard.

Because presupposition is the core motivation for deflationism, §3.1 begins with some crucial but unappreciated facts about presupposition. §3.2 then introduces deflationism. I'll claim that all moral realists should accept deflationism because it's the uniquely correct way for a Kratzer-style approach to incorporate the property *being-the-moral-standard*.

In particular, I claim that that property must be part of a presupposed, not-at-issue commitment associated with the moral use of modals, and that deflationism follows from that claim. Deflationism has been defended at length elsewhere (Perl 2020b, ms; Perl and Schroeder 2019); that work in effect gives the semantic-cum-pragmatic underpinnings of deflationism, while this paper describes one of its central philosophical applications.

3.1 Presupposition: background

Philosophers in general have missed an important quirk in our concept of knowledge, which arises from commitments that are presupposed. My ambition in this section is to convince you that there is an important quirk that has been largely missed.

Presuppositions are commitments of an utterance that are interpreted as backgrounded and not the main point. Consider the sentence “it wasn’t Bill who ate the cookies”. Uses of that sentence are associated with the commitment that *someone* ate the cookies – if you sincerely utter it, competent hearers will infer that you accept that commitment. But competent hearers will also recognize that the main point of the utterance is elsewhere. The main point is to convey something about Bill’s inactivity, rather than the current state of the cookies. In general, presuppositions are interpreted as not the main point; following Craige Roberts (2012), I’ll say that they’re *not-at-issue* commitments.

Presuppositions give rise to an interesting but unappreciated species of knowledge. Suppose that I suspect that someone did eat the cookies. Suppose further that I know that Bill didn’t eat the cookies – say, because I’ve been watching him all day, and know that he didn’t go anywhere near them. Is (*) then true?

(*) I know that it wasn’t Bill who ate the cookies.

Yes! (*) is true, given those suppositions. That is, knowing the at-issue content and merely suspecting the not-at-issue content can be enough.

(*) illustrates an important general lesson: that we can use knowledge reports to assert something true given knowledge of the at-issue commitment (that Bill didn’t eat the cookies) and mere acceptance of the presupposed, not-at-issue commitment (that someone did). (After all, suspecting something can sometimes be enough for temporarily accepting it.) Stalnaker suggested in general that *acceptance* is the attitude appropriate for not-at-issue commitments. And this kind

of point is uncontroversial among linguists – Heim (1992) is a *locus classicus*, though much work builds on hers.¹⁰ If you find this point surprising, your surprise is evidence that philosophers are in general unaware of an important way knowledge reports work.

The general lesson is a lesson specifically about the distinction between at-issue and not-at-issue commitments. Suppose that I merely suspect that Bill didn't eat the cookies, but I know that someone did. Could (*) then be true?

(*) I know that it wasn't Bill who ate the cookies.

No, it couldn't. In order for (*) to be true, I do have to know the at-issue content that Bill didn't eat the cookies. The felicity of (*) under the earlier suppositions shows something specific about not-at-issue commitments, rather than something general about knowledge.

3.2 Deflationism, what

The rough statement of deflationism is that knowledge of an ordering can ground moral knowledge for someone who *presupposes* that that ordering is the moral standard. This section develops that rough statement into something more precise.

There is some ordering that best fits my dispositional moral beliefs, and my dispositions about it allow me to know about it. I'll call that ordering \mathbf{c}_1 . I use a singular term (' \mathbf{c}_1 ') to refer to my own ordering, to emphasize that the attitudes are not attitudes about myself. They are rather about a particular ordering \mathbf{c}_1 – the fact that \mathbf{c}_1 is the ordering that best fits my moral beliefs figures only in the *metasemantic* explanation of why my attitudes are about that ordering, not in the attitudes' content. Since that ordering is the one that best fits my dispositional moral beliefs, I can know about it, and indeed know about it *a priori*. That's the kind of knowledge that §2.2 emphasized.

Deflationism holds that knowledge of my ordering partially grounds moral knowledge. But it's only a *partial* ground; it grounds moral knowledge only in combination with a further state: the state of *accepting* that the relevant part of that ordering lines up with the moral standard. For example, knowing that killing is usually wrong can be grounded in the following two states about my ordering \mathbf{c}_1 :

¹⁰Examples include Paul Dekker (2008), Lauri Karttunen and Stanley Peters (1979), Robert van Rooij (2005, 2010), David Oshima (2006), and Yasutada Sudo (2012); my (2020c) explains why any viable account of presupposition triggers must vindicate the observations described in the main text.

- knowledge that killing is usually wrong_{c₁}, plus
- accepting that the relevant part of **c₁** lines up with the moral standard

These two states are ones we'd expect to be crucial given a focus on presupposition: acceptance is that attitude appropriate for presupposed, not-at-issue commitments.

To accept a proposition is to treat it as true. And acceptance is not normed by knowledge, in the way that belief plausibly is. The only general norm on acceptance is coherence, both logical and probabilistic.¹¹ You can appropriately criticize what someone accepts if it's inconsistent with other things that they accept or believe. It's inappropriate to criticize them for other reasons.

Deflationism holds that two crucial mental states can ground knowledge that killing is usually wrong.

- knowledge that killing is usually wrong_{c₁},
- accepting that the relevant part of **c₁** lines up with the moral standard

Crucially, though, these two states don't always ground moral knowledge. The rest of the section will introduce the complications where they don't ground moral knowledge, so that I can later explain why those complications are irrelevant for our attributions of *apriori* knowledge. But I encourage readers to skip ahead to §4 whenever they feel that they have a good enough grasp of deflationism. Even at this point, you might already see enough to see how the basic vindication of *apriori* knowledge will go.

The first complication is that knowledge attributions obey an additional constraint. The additional constraint is that an attributor will attribute moral knowledge to an addressee only when the attributor herself agrees with the addressee about the relevant part of the moral standard. For example, you'll take me to know that killings are usually wrong only if you also accept that the relevant part of **c₁** lines up with the moral standard. And you accept that it does if you agree with me about the considerations that make killing wrong: if you agree that killings are wrong unless they're in self-defense, and even then they're wrong if there was a non-lethal alternative, and so on. I'll call this final feature of deflationism its **DEFLATIONIST FACTIVITY**: an attributor attributes moral knowledge only when she also accepts that the addressee's ordering matches the relevant part

¹¹See Stalnaker (2002) for a full account of this attitude.

of the moral standard. (DEFLATIONIST FACTIVITY faces a range of familiar formal problems; my (2020a) gives my solution.)

DEFLATIONIST FACTIVITY guarantees that deflationism avoids the problems from §1 for extreme versions of a Jackson-style view. I objected that those kinds of views predict that (*) is true:

(*) Rossians know that promise-breaking is sometimes wrong just because it's promise-breaking, while utilitarians know that promise-breaking is never wrong just because it's promise-breaking.

As noted earlier, extreme Jackson-style views take Rossians to be thinking about something subtly different than utilitarians are. Rossians are thinking about $wrong_R$, and utilitarians about $wrong_U$. Since they're thinking about different things, they can both know what they believe. Deflationism *agrees* with Jackson-style views that Rossians and utilitarians are thinking about different things.

Despite an element of agreement with Jackson-style views, deflationism nonetheless predicts that (*) is *always false*. Given DEFLATIONIST FACTIVITY, (*) is true only if the relevant parts of the Rossian and the utilitarian moral orderings *both* line up with the moral standard. But it's impossible for the two orderings to both line up with the moral standard, because they disagree about promise-keeping. So the crucial difference between deflationism and Jackson-style views is that Jackson-style views takes moral knowledge to be *fully* grounded in knowledge of something analytic. In contrast, deflationism holds that moral knowledge is only *partially* grounded in knowledge of something analytic – and the other ground is what guarantees that (*) is always false.

DEFLATIONIST FACTIVITY differs from CLASSICAL FACTIVITY:

DEFLATIONIST FACTIVITY: an attributor attributes moral knowledge only when she also accepts that the attributee's ordering matches the relevant part of the moral standard.

CLASSICAL FACTIVITY: someone has moral knowledge only when it's **true** that the attributee's ordering matches the relevant part of the moral standard.

Given DEFLATIONIST FACTIVITY, attributions of moral knowledge are all attributions *from within a 'perspective'* – from the perspective of a person who herself accepts substantive propositions about the moral standard. In fact, you and I might take different knowledge attributions to be correct because we accept different propositions

about the moral standard. DEFLATIONIST FACTIVITY thus resembles the sophisticated sort of expressivism that insists that (*) is always false because – very roughly – it’s false from every ‘perspective’. Unlike expressivism, though, deflationism gives a purely realist characterization of the perspectives, individuating them by the propositions accepted. Classical Factivity holds of God’s attributions of moral knowledge, but need not hold for you or me.

The other complication about deflationism is that moral knowledge is context-sensitive in a new way. We attribute knowledge very differently in what I’ll call *Easy Contexts* than we do in what I’ll call *Hard Contexts*. What I’ve said so far only describes knowledge in Easy Contexts.

In Hard Contexts, moral knowledge does require knowledge and not mere acceptance about the moral standard – about the fundamental moral facts that moral realists posit. In such contexts, moral knowledge might require knowing that c_1 is the moral standard, or it might require knowing that according to the moral standard, killing is wrong. Hard Contexts are thus contexts where moral knowledge works just like you’ve always expected moral knowledge to work if moral realism is true. The deflationist’s key innovation is to introduce Easy Contexts; you’re already familiar with Hard Contexts.

The final key feature of deflationism is its account of shifts between Easy Contexts and Hard Contexts. *The only way to shift someone into a Hard Context from an Easy Context is to get them to stop accepting substantive propositions about the moral standard.* I’ll call this thesis the SHIFTING THESIS. The Shifting Thesis is one of the core differences between deflationism and more familiar kinds of epistemic contextualism. For example, Lewis (1996) incorporates a ‘Rule of Attention’ to determine shifts into higher-standards contexts. In the present setting, a Rule of Attention would mean that I’m *automatically* in a Hard Context if I’m paying attention to the possibility that c_1 (my current ordering) might not be the moral standard. Deflationism does *not* incorporate a Rule of Attention or anything similar. Getting me to stop accepting substantive propositions is the *only* way to shift me into a Hard Context. Now a skeptic who wants to shift me into a Hard Context may well *start* by getting me to pay attention to the possibility that c_1 isn’t the moral standard. But such a skeptic needs to do a great deal more – whereas on a Lewis-style view, such a skeptic succeeds automatically when she brings the possibility to attention.

Deflationism is an account of *all* epistemic properties, not just knowledge. It is, for example, also an account of justification. In

Easy Contexts, I can be *justified* in believing that killing is usually wrong if I'm justified in believing that killing is usually wrong_{c₁} while also merely accepting that **c₁** is the moral standard. And I can know *a priori* that killing is usually wrong if I know *a priori* that killing is usually wrong_{c₁} while also accepting that **c₁** is the moral standard. Though I'll focus on moral knowledge, the focus is merely expository.

3.3 Deflationism, why

Suppose that moral uses of 'killing is usually wrong' assert that killing is usually wrong_{c₁}, while presupposing that **c₁** is the moral standard. Deflationism would then be true. Knowledge that killing is usually wrong could then be grounded in:

- knowledge that killing is usually wrong_{c₁}, plus
- accepting that (the relevant part of) **c₁** is the moral standard

Knowledge can be grounded in these two states because acceptance of the presupposed, not-at-issue commitment can combine with knowledge of the at-issue commitment to ground knowledge, as §3.1 emphasized.

The other distinctive features of deflationism also follow immediately from this supposition about what's presupposed.

- DEFLATIONIST FACTIVITY follows because 'knows' is a hole for presupposition – that is, a knowledge attribution is appropriately assertable only if the speaker also accepts the presupposition.
- Hard Contexts exist because presuppositions can be locally accommodated under attitude reports – that is, attitude reports can be interpreted as targeting the presupposed, not-at-issue commitment as well as the at-issue commitment. ('I know that it wasn't Bill who ate the cookies' can be used to communicate that I know the not-at-issue commitment that someone ate them as well as the at-issue commitment that Bill didn't.)
- Changing what someone accepts is the only way to shift them into a Hard Context because local accommodation only happens given changes to what's accepted.

It's these claims that are defended elsewhere (Perl 2020b, ms; Perl and Schroeder 2019).

In that other work, I also develop what I take to be decisive evidence that the property *being-the-moral-standard* is part of a presupposed,

not-at-issue commitment. The evidence is that so supposing is the only way to preserve Kratzer’s elegant explanation of the duality of modals once we integrate the realist property *being-the-moral-standard*. More generally, I take the distinction between what’s at-issue and what’s not- to be a distinction in thought as well as in talk. So I take my evidence about what’s not-at-issue to be evidence about our *concept* of moral knowledge as well as our *talk* about moral knowledge.

I close with some quick technical points, but I encourage most readers to skip ahead to §4. I opened by supposing that moral utterances are associated with the presupposed, not-at-issue commitment that c_1 is the moral standard. That supposition illustrates one way that deflationism could be true. More generally, though, deflationism is true if the property *being-the-moral-standard* is *part* of a presupposed, not-at-issue commitment. So it’d also be true if a *subpropositional* commitment is interpreted as presupposed and not-at-issue. For instance, a propositional function from an object/ discourse referent to the proposition that it’s the moral standard could also be interpreted as presupposed and not-at-issue.

Deflationism is true as long as the fundamental realist property is interpreted as presupposed and not-at-issue, however it otherwise figures in what’s communicated. Suppose, for instance, that ‘killing is usually wrong’ communicates the existential proposition that $[\exists x: x \text{ is the moral standard}]$ (x usually forbids killing), with ordinary uses presupposing that x is the moral standard. Then ‘I know that killing is wrong’ would communicate:

(2) I know that $[\exists x: x \text{ is the moral standard}]$ (x usually forbids killing)

Consider the benchmark account of presupposed commitments that Irene Heim (1992) developed. Roughly, she suggests *A knows that S* is true iff the propositions that A knows and the presuppositions of S that A accepts jointly entail that S.¹² The deflationist’s two mental states would then ground moral knowledge: the proposition that I know (that killing is usually wrong _{c_1}) plus the presupposed, not-at-issue commitments that I merely accept (that c_1 is the moral

¹²A complication is that Heim herself formulates her account by appeal to belief rather than acceptance: what matters for her is the propositions that I know and the propositions that I *believe*, rather than the propositions that I accept. But I take that formulation to be a simplification so that she can focus on the semantic questions that really concern her; I assume that Stalnaker (2002) is right that it’s really acceptance rather than belief that matters. That’s one lesson of the earlier example with Bill and the cookies.

standard) jointly entail (2)'s complement. Though I initially formulated the presuppositional thesis with singular propositions, I can also formulate it given any realist account of what moral utterances communicate.

4 Vindicating the hallmarks of moral knowledge

Deflationism vindicates the *aprioricity* of moral knowledge. Consider my confidence that I know *a priori* that I ought to keep my promises. Since I'm confident that I ought to, I do accept that the part of my ordering about promise-keeping is the moral standard. As a result, I'm in an Easy Context, and I have one of the two mental states that ground moral knowledge there. I also accept that DEFLATIONIST FACTIVITY is satisfied, because I do accept that my ordering is the moral standard.

And as §2.2 emphasized, I also know *a priori* that my own ordering requires promising-keeping, given my disposition see promise-keeping as what I ought_{c₁} to do. So I also attribute to myself the other state that can ground *a priori* moral knowledge. (Asking whether what I accept is *a priori* doesn't make sense; *aprioricity* in the sense that interests us is a property of epistemic states, like knowledge or justification: it tells us *how* the epistemic properties are available.)

We can now return to a question that §2 left open. It seems like I can know *a priori* about an ordering, and just be lucky that that ordering is the moral standard. §2 closed by asking whether *that* kind of luck excludes knowledge. We're finally in a position to answer. Deflationism gives a distinctive answer: that that luck is compatible with *a priori* moral knowledge, because that luck is luck in what I *accept*, rather than luck in what I *know*. There's no luck in knowing that I ought_{c₁} to keep my promises, because c₁ is the ordering, whichever it is, that best fits my dispositions. Deflationism gives a constitutive explanation of *that* knowledge. And if there's any luck here, the luck would be *capacity* epistemic luck, which doesn't exclude knowledge, as §2 noted.

In contrast, I *am* lucky in accepting something true. But *that* luck is irrelevant: knowledge is not the norm on acceptance, so anti-luck constraints on knowledge simply won't apply. Moral knowledge is compatible with a kind of luck because one of the states that grounds moral knowledge in Easy Contexts is compatible with that kind of luck. If deflationism is true, there seems to be a new and unappreciated kind of innocuous epistemic luck that's compatible with knowledge: luck in what's accepted.

At this point, you might object that certain kinds of acceptance

just can't ground genuine moral knowledge. For instance, I might accept that c_1 is the moral standard because I used a series of coin-flips to decide what to accept as the moral standard. *That* kind of acceptance doesn't ground moral knowledge!

I agree: only certain kinds of acceptance ground genuine moral knowledge. To introduce the kind of acceptance that I take to matter, suppose you and I confront a common lunch bill. It includes neither tax nor tip. We calculate different answers about the tax due, and independently calculate different answers about the tip due. (Suppose that they're independent questions: we're calculating both from the base bill.) Neither of us is very confident in each answer: we're just .6 that each answer is correct. I say that the total is \$64.38, and you say that the total is \$61.25. Our credence in each answer is just .36. I nonetheless accept that the total is \$64.38 because my credence in that proposition is higher than any of the relevant alternatives. It's higher than my credence that the total is \$61.25, or that it's \$60.25, or ... I say that this example illustrates the attitude of *ur-acceptance*, where I ur-accept that p iff my credence that p is higher than my credence in any relevant alternative.

Ur-acceptance is the species of acceptance that matters in ordinary moral contexts. We attribute moral knowledge only when someone ur-accepts that the relevant ordering is the moral standard. That's why we'll refuse to attribute moral knowledge when they accept propositions about the moral standard because of coin-flips. Ur-acceptance is what matters in ordinary moral contexts because of our *practical* interests in those contexts. For instance, we're often interested in what we can rely on others to do, or at least what they think they should do. We can rely on what someone ur-accepts in a way we can't rely on what they accept as a result of coin-flips. I explain why ur-acceptance is the species of acceptance that matters in ordinary contexts in more detail in Chapter 5 of my (ms).

Crucially, though, coherence remains the only general norm on ur-acceptance. Since I can't know p given a credence just of .36 in p , knowledge can't be a norm on ur-acceptance. The philosophical upshots of deflationism remain intact even when we focus just on what's ur-accepted, since the crucial point about acceptance was that coherence is its norm. I'll continue to talk about acceptance in what follows, but I mean to restrict attention just to ur-acceptance.

What I've done so far is explain why deflationism explains the *aprioricity* of moral knowledge in Easy Contexts. I now claim that Easy Contexts are the only ones that matter for the evaluation of moral naturalism. They're the only ones that matter because of

the way that non-naturalists reason about the *aprioricity* of moral knowledge. When a non-naturalist argues that moral knowledge is available *apriori*, she's attributing a piece of moral knowledge to herself, and thinking about what it would take for that attribution to be true. But *that very attribution* reveals that she is in an Easy Contexts. Hard Contexts are contexts where the speaker *doesn't* accept that relevant part of the moral ordering lines up with the moral standard. So they wouldn't be contexts where someone would attribute knowledge to herself. As a result, contexts where someone takes herself to have *apriori* knowledge must be Easy Contexts.

Moreover, deflationism gives a new diagnosis of why the non-naturalists' arguments are misleadingly plausible. The diagnosis is that there are contexts where their arguments *are* plausible: Hard Contexts. Non-naturalists are implicitly imagining what the naturalist would say in those contexts. Consider, for example, a moral skeptic who insists that she has no grounds for accepting substantive propositions about the moral standard. She's in a Hard Context because she doesn't accept any such substantive propositions. In giving reasons for the skeptic to change her mind, the naturalist might be appealing to facts known only *a posteriori* – facts about whatever natural properties the naturalist takes to be fundamental. The evidence given wouldn't be available *apriori*.

Even though moral truths aren't available *apriori* in Hard Contexts, they're still available *apriori* in Easy Contexts. But as just noted, the *aprioricity* hallmark of moral knowledge is well-motivated only in Easy Contexts, because our confident conviction in that hallmark rests on our accepting substantive propositions about the moral standard. Naturalists can smoothly explain that hallmark there. But the non-naturalists is illegitimately imagining what happens in Hard Contexts – which is altogether irrelevant for the hallmark that motivates the non-naturalist in the first place. Non-naturalists mistakenly take the *aprioricity* of moral knowledge to be evidence against moral naturalism because they're illegitimately switching between Easy and Hard Contexts.

More generally, my deflationist proposal dissolves a range of objections to moral naturalism, for similar reasons.

4.1 Minimalist epistemology

Rawls articulated another hallmark of moral knowledge:

we do not require of a moral or political judgment that the reasons for it be related to an appropriate causal process, or require an explanation of it within cognitive

psychology. It is enough that the reasons offered [in support of some normative judgment] be sufficiently strong. We explain our judgment, so far as we do, simply by going over the grounds for it: the explanation lies in the reasons we sincerely affirm. What more is there to say except to question our sincerity and reasonableness? (Rawls 2005, 118)¹³

Rawls is here suggesting that we justify some claims to moral knowledge by describing the reasons that favor the proposition that we're claiming to know. Once we've done that, we don't need to do anything else.

Knowledge of natural facts seems different. I might defend a claim to knowledge that something is green by giving my reason: that it seems green to me. But this defense only takes me so far. The reason I'm giving is genuine only if it's related to "an appropriate causal process": if my seeming is the upshot of a reliable process. Rawls is suggesting that the moral case is different. In the moral case, we need only have true beliefs about the reasons that bear on the question, and how those reasons weigh together. Our beliefs don't need to be the upshot of an appropriate causal process. I'll say that Rawls is describing a *minimalist* epistemology for moral claims: minimalist in the sense that moral knowledge doesn't need to be backed by reliable processes.

There is fairly wide agreement on the first two hallmarks of moral knowledge: its apriority and its autonomy from scientific inquiry. There is much less agreement about whether moral knowledge has the minimalist epistemology that Rawls suggests.¹⁴ But I think many philosophers do agree with Rawls, and take the minimalist epistemology as powerful independent evidence against moral naturalism.¹⁵

¹³Scanlon is similar: "as in the previous case, the process here is to try to characterize the potential reason more fully, to ask whether it seems, so characterized, to be a relevant reason for the attitude in question. In addition, one can look for other cases on which it would have a bearing if it were a good reason, to see whether it seems to be a reason in those cases, to test one's reaction in those cases for signs of unreliability, to consider the plausibility of alternative explanations of these reactions, and so on" (Scanlon 1998, 68).

¹⁴See Barry Maguire (2014) and Ralph Wedgwood (2016) for some representative concerns, articulated as concerns about this kind of minimalist epistemology as it appears in Scanlon (2014).

¹⁵I take a commitment to this sort of minimalist epistemology to be tacitly important for a wide range of non-naturalists: philosophers like G. E. Moore (1903), W. D. Ross (1930), Thomas Nagel (1986), Ronald Dworkin (1996), Jean Hampton (1998), Derek Parfit (2011), and T. M. Scanlon (2014). I've encountered people who claim that these philosophers adopt minimalist epistemology in order to make their non-naturalism work, rather than seeing this minimalist

Since this hallmark has been so important, it's worth discussing with the others, even though it's more controversial.

Deflationism will allow naturalists to capture this minimalist epistemology in Easy Contexts. In those contexts, the knowledge required is knowledge about a particular moral ordering. And that knowledge is grounded in my dispositions. So it won't require backing by any 'appropriate causal process', any more than our logical knowledge needs backing by an appropriate causal process. It's again helpful to situate Rawls' point in to the earlier discussion of luck. If moral knowledge rested on knowledge of natural properties, our moral beliefs would be luckily true unless backed by a causal process. But since moral knowledge rests on mere acceptance about natural properties, we can have it even in the absence of a causal process that excludes luck. That's why it seems to us like moral epistemology is as minimalist as Rawls and others suggest.

4.2 The autonomy of moral inquiry

Another challenge to moral naturalists appeals to the *autonomy* of moral inquiry from scientific inquiry. Moral inquiry isn't fully autonomous from scientific inquiry, since scientific investigation may illuminate some moral questions. For example, it may illuminate our obligations not to pollute. But it doesn't seem like scientific inquiry *settles* moral questions. As Russ Shafer-Landau puts the point, "ethics is an autonomous domain of inquiry; ethical concerns are not reducible without remainder to any other kind of concern, and ethical investigation is not properly conceived of as a subset of any other recognized discipline" (Shafer-Landau 2003, 52).

It is hard for naturalists to explain why moral inquiry would be autonomous from scientific inquiry. If moral naturalism is true, it seems like moral inquiry should just be a species of scientific inquiry. Scientific inquiry seems like the appropriate way to investigate natural facts, and the naturalist holds that moral facts are natural facts.

Deflationism gives a distinctively *coherentist* account of moral inquiry; we justify our moral beliefs by showing that they cohere with our other moral beliefs. My belief that killing is usually wrong is justified in part in the minimalist way that §4.1 described – by recognizing the factors that matter according to c_1 . In addition, I'll take myself to be justified only because I accept that c_1 is the moral standard. I expect that most arguments about whether I'm justified

epistemology as independent evidence against moral naturalism. But that claim is simply mistaken for at least some philosophers – particularly for Rawls and Scanlon, as the earlier quotes illustrate.

are really arguments about whether c_1 is the moral standard. So my moral justification is distinctively grounded in what I accept. And coherence is the only general norm on coherence.

Now many readers will immediately object that coherence doesn't itself justify: if I form beliefs at random about my physical environment, those beliefs won't be justified no matter how coherent they are. I'm happy to agree. Coherence itself doesn't justify empirical beliefs. But deflationism is *not* a view about empirical beliefs – it takes empirical beliefs to work just like you've always assumed. If deflationism is true, moral justification differs radically from empirical justification. What I *accept* grounds moral justification, even though it doesn't ground empirical justification. Coherentism about moral justification doesn't require coherentism about empirical justification.

My moral beliefs are justified to the extent that they cohere with my other moral beliefs. That's why we see our moral intuitions as evidence about moral reality *independently* of whether we see them as a scientific or proto-scientific way of discovering moral facts. Coherence within our pure moral beliefs justifies our pure moral beliefs. More generally, deflationism doesn't take our moral intuitions to be direct intuitions about natural properties; they're instead intuitions about our individual moral orderings, which are accepted to have the natural property *being-the-moral-standard*. (And coherence is the norm on acceptance.) Of course our intuitions don't need to be explained as a scientific or proto-scientific method.

5 Wrapping up

I've shown how deflationism allows moral naturalists to capture three central hallmarks of moral knowledge: its *aprioricity*, its *autonomy from scientific inquiry*, and its *minimalist epistemology*. Now you might doubt that some of these hallmarks are genuine, or you might think that naturalists have other options for explaining them. But if deflationism is true, there's no need to contest these hallmarks, and no need to work at working out a different explanation of them: deflationism takes care of it all. And I think there is *decisive* evidence that deflationism is true. I haven't given that evidence here, since it's given elsewhere (Perl 2020b, ms; Perl and Schroeder 2019) – I've rather focused on describing the upshots of deflationism for debates about moral naturalism.

Deflationism frees us to focus on the debates in moral *metaphysics* that should have been central all along.¹⁶ Moral naturalists do

¹⁶Of course, instances of Frege's puzzle will also remain: someone can

need to identify natural properties that are plausible candidates for moral properties. Identifying those properties is an exercise in first-order normative ethics; hedonistic act utilitarianism is a classic if implausible option. Deflationism frees moral naturalists to focus our efforts there, by dissolving the epistemic objections that have distracted us.

References

- Boyd, Richard. 1988. "How to be a Moral Realist." In Geoffrey Sayre-McCord (ed.), *Essays on Moral Realism*, 181–228. Cornell University Press.
- Brandt, Richard. 1979. *A Theory of the Good and the Right*. Oxford: Oxford University Press.
- Brink, David. 1989. *Moral Realism and the Foundations of Ethics*. Cambridge: Cambridge University Press.
- . 2001. "Realism, naturalism, and moral semantics." *Social Philosophy and Policy* 18:154–76.
- Brown, Campbell. 2011. "A New and Improved Supervenience Argument for Ethical Descriptivism." In Russ Shafer-Landau (ed.), *Oxford Studies in Metaethics* 6, 205–18. Oxford: Oxford University Press.
- Copp, David. 2000. "Milk, Honey, and the Good Life on Moral Twin Earth." *Synthese* 124:113–137.
- . 2007. *Morality in a Natural World: Selected Essays in Metaethics*. New York: Cambridge University Press.
- Dekker, Paul. 2008. "A multi-dimensional treatment of quantification in extraordinary English." *Linguistics and Philosophy* 31:101–127.
- Dowell, Janice. 2013. "Naturalism, Ethical."
- . 2016. "The Metaethical Insignificance of Moral Twin Earth." *Oxford Studies in Metaethics* 11:1–26.
- Dreier, James. 1992. "The Supervenience Argument against Moral Realism." *Southern Journal of Philosophy* 30(3):13–38.

competently deny that the correct naturalist reduction captures the property *being-the-moral-standard*, and we need to explain why; Nicholas Laskowski (2019) gives one possible approach.

- Dunaway, Billy. 2015. "Supervenience Arguments and Normative Non-naturalism." *Philosophy and Phenomenological Research* 91 (3):627–655.
- Dunaway, Billy and McPherson, Tristram. 2016. "Reference Magnetism as a Solution to the Moral Twin Earth Problem." *Ergo* 3.
- Dworkin, Ronald. 1996. "Objectivity and truth: you'd better believe it." *Philosophy and Public Affairs* 25:87–139.
- Foot, Philippa. 2001. *Natural Goodness*. Oxford: Oxford University Press.
- Hampton, Jean. 1998. *The Authority of Reason*. Cambridge: Cambridge University Press.
- Heim, Irene. 1992. "Presupposition Projection and the Semantics of Attitude Verbs." *Journal of Semantics* 9 (3):183–221.
- Huemer, Michael. 2005. *Ethical Intuitionism*. Houndmills: Palgrave MacMillan.
- Hursthouse, Rosalind. 1999. *On Virtue Ethics*. Oxford: Clarendon Press.
- Jackson, Frank. 1998. *From Metaphysics to Ethics: A Defence of Conceptual Analysis*. Oxford: Clarendon Press.
- Karttunen, Lauri and Peters, Stanley. 1979. "Conventional implicature." In Choon-Kyu Oh and David Dinneen (eds.), *Syntax and semantics vol. 11: Presupposition*, 1–56. New York: Academic Press.
- Kolodny, Niko and MacFarlane, John. 2010. "Ifs and Oughts." *Journal of Philosophy* 107 (3):115–143.
- Kratzer, Angelika. 1977. "What 'Must' and 'Can' Must and Can Mean." *Linguistics and Philosophy* 1:337–355.
- . 1981. "The Notional Category of Modality." In H J Eikmeyer and H Rieser (eds.), *Words, Worlds, and Contexts*, 38–74. New York: de Gruyter.
- . 2012. *Modals and Conditionals. New and Revised Perspectives*. Oxford: Oxford University Press.

- Laskowski, Nicholas. 2019. "The sense of incredibility in ethics." *Philosophical Studies* 176 (1):93–115.
- Laskowski, Nicholas and Finlay, Stephen. 2017. "Conceptual Analysis in Metaethics." In Tristram McPherson and David Plunkett (eds.), *Routledge Handbook of Metaethics*. Routledge.
- Laskowski, Nicholas and Howard, Nathan Robert. 2019. "The world is not enough." *Nous* doi:10.1111/nous.12293.
- Lewis, David. 1989. "Dispositional Theories of Value." *Aristotelian Society, Supplementary Volume* 63:113–137.
- . 1996. "Elusive Knowledge." *Australasian Journal of Philosophy* 74:549–567.
- Maguire, Barry. 2014. "T. M. Scanlon Being Realistic About Reasons." *Times Literary Supplement* 26.
- McPherson, Tristram. 2015. "What is at Stake in Debates among Normative Realists?" *Nous* 49 (1).
- Moore, G E. 1903. *Principia Ethica*. Cambridge: Cambridge University Press.
- Nagel, Thomas. 1986. *The View from Nowhere*. Oxford: Oxford University Press.
- Oshima, David. 2006. *Perspectives in Reported Discourse*. Ph.D. thesis, Stanford University.
- Parfit, Derek. 2011. *On What Matters*, volume 2. Oxford: Oxford University Press.
- Perl, Caleb. ms. *Deflating moral epistemology*.
- . 2020a. "A user's guide to hybrid tools." *Mind* 129 (513):129–158.
- . 2020b. "Might moral epistemologists be asking the wrong questions?" *Philosophy and Phenomenological Research* 100 (3):556–585.
- . 2020c. "Presuppositions, attitudes, and why they matter." *Australasian Journal of Philosophy* 98 (2):363–381.
- Perl, Caleb and Schroeder, Mark. 2019. "Attributing error without taking a stand." *Philosophical Studies* 176:1453–1471.

- Pigden, Charles. 2012. "Identifying Goodness." *Australasian Journal of Philosophy* 90(1):93–109.
- Pritchard, Duncan. 2005. *Epistemic luck*. Oxford University Press: Oxford.
- Railton, Peter. 1986. "Moral Realism." *The Philosophical Review* 95 (2):163–207.
- Rawls, John. 1975. "The Independence of Moral Theory." *Proceedings and Addresses of the American Philosophical Association* 48:5–22.
- . 2005. *Political Liberalism*. New York: Columbia University Press.
- Ridge, Michael. 2007. "Anti-Reductionism and Supervenience." *Journal of Moral Philosophy* 4:330–48.
- Roberts, Craige. 2012. "Information Structure: Towards an integrated formal theory of pragmatics." *Semantics and Pragmatics* 5 (6):1–69.
- Ross, W D. 1930. *The Right and the Good*. Oxford: Oxford University Press.
- Scanlon, T M. 1998. *What We Owe to Each Other*. Cambridge, MA: Harvard University Press.
- . 2014. *Being Realistic about Reasons*. New York: Oxford University Press.
- Schroeder, Mark. 2005. "Realism and Reduction: The Quest for Robustness." *The Philosophers' Imprint* 5 (1).
- . 2007. *Slaves of the Passions*. Oxford: Oxford University Press.
- . 2019. "A common subject for ethics." .
- Shafer-Landau, Russ. 2003. *Moral Realism: A Defence*. Oxford: Oxford University Press.
- . 2006. "Ethics as Philosophy – A Defence of Ethical Nonnaturalism." In Russ Shafer-Landau (ed.), *Metaethics after Moore*, 209–232. Oxford University Press.
- Stalnaker, Robert. 2002. "Common Ground." *Linguistics and Philosophy* 25 (5-6):701–721.

- Sturgeon, Nicholas. 2006. "Ethical Naturalism." In David Copp (ed.), *The Oxford Handbook of Ethical Theory*, 91–121. New York: Oxford University Press.
- Sudo, Yasutada. 2012. *On the Semantics of Phi Features on Pronouns*. Ph.D. thesis, Massachusetts Institute of Technology.
- Thompson, Michael. 2012. *Life and Action*. Cambridge, MA: Harvard University Press.
- van Rooij, Robert. 2005. "A modal analysis of presupposition and modal subordination." *Journal of Semantics* 22:281–305. ISSN 0167-5133.
- . 2010. "Presupposition: An (un)common attitude?" In Rainer Bauerle, Uwe Reyle, and Thomas Zimmermann (eds.), *Presuppositions and Discourse: Essays Offered to Hans Kamp*. Emerald Group Publishing Limited.
- Viebahn, Emanuel and Vetter, Barbara. 2016. "How many meanings for 'may'? The case for modal polysemy No Title." *Philosopher's Imprint* 16 (10):1–26.
- Wedgwood, Ralph. 2007. *The Nature of Normativity*. Oxford: Oxford University Press.
- . 2016. "Review: Being Realistic about Reasons." *The Philosophical Review* 66 (262):213–217.
- Willer, Malte. 2014. "Dynamic Thoughts on Ifs and Oughts." *Philosophers' Imprint* 14 (28):1–30.
- Yablo, Stephen. 2000. "Red, Bitter, Best." *Philosophical Books* 41:13–23.