

Might moral epistemologists be asking the wrong questions?

This paper suggests that moral epistemologists have been focusing on the wrong questions, or at least ignoring some of the right ones. They've been ignoring some of the right questions because they've made a mistake about the semantics and pragmatics of moral discourse.

For concreteness, I'll focus on moral skeptics, who argue that none of us have any substantive moral knowledge. Different skeptics offer different arguments. Some appeal to general empiricist strictures on what knowledge requires; others to pervasive moral disagreement; others to the dispensability of moral facts in scientific explanation; others to other considerations.¹ Moral skeptics usually focus on their attention on moral realists, who take moral properties to be independent of and more fundamental than our individual evaluative attitudes. Skeptics want to show that we don't know what's permitted or required if moral realism is true. (They tend to be much less interested in showing that moral anti-realism (a la A. J. Ayer (1936), say) has skeptical consequences.) In discussing arguments for moral skepticism, it thus makes sense to assume that some kind of moral realism is true, to focus on the issues that animate most moral skeptics.

If this paper is right, moral skeptics need to address a very new question before they can press their challenge. Moreover, the traditional arguments for moral skepticism do not illuminate the new question. So central questions in moral epistemology would look very different if this paper is correct. I myself suspect there is no way for moral skeptics to press their challenge, once we fully internalize the lessons from this paper. Unfortunately, though, there won't be space here to substantiate this suspicion. My goal will rather be to convince you that we need to think about moral epistemology in a very new way.

1 Introducing the presuppositional conjecture

This paper will link questions about moral discourse to questions about moral epistemology. I'll begin in this section by focusing on the standard semantics for modals that Angelika Kratzer developed, and describing the different ways that moral realists can accept that semantics.

¹David Hume (1739) leveraged empiricist strictures in arguing that moral rationalism leads to skepticism. Sextus Empiricus (2000) argued from disagreement; Katia Vavova (2014) and Kieran Setiya (2012) both have illuminating discussions of arguments from disagreement. Gilbert Harman (1977) appeals to the dispensability of moral facts in scientific explanation, and Sharon Street (2006) appeals to scientific explanations of moral belief. For more detail about these arguments and others, see Walter Sinnott-Armstrong (2006).

1.1 The standard account of moral discourse

We use modal verbs like *can*, *must*, and *have to* and modal adjectives like *wrong* and *permissible* in a wide range of ways. We use them to talk about the demands of morality, or the demands of prudence, or 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.² We should expect the best semantics for modal terms to explain the range of uses.

Angelika Kratzer (1977, 1981, 2012) has developed the standard account. She suggests that modal terms 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 an argument that combines with the semantics for the modal term to determine a complete, truth-evaluable proposition.

Kratzer's contextualist semantics crisply explains why modal terms have a range of different uses in natural language. Occurrence of "have to" make the same semantic contribution to different uses of (1). But (1) can still be used to communicate different propositions, because the context can supply different arguments to saturate the semantic contribution of "have to". This contextualist account unifies the different uses of modal terms. The contextualist predicts that, as soon as we understand one use of (1), we can understand the other uses too, because understanding one use of "have to" involves understanding the semantic contribution of "have to", and "have to" makes the very same semantic contribution in each of the other uses. And that's why modal terms have a robust range of shared uses across different natural languages: when you understand one of the uses, you're in a position to understand the rest.

Now the details of Kratzer's view are very complicated, in ways that will be irrelevant in this paper. We can abstract from them by saying that the context supplies an *ordering* of worlds that ranks possible worlds in relation to each other.³ A modal like "have to" takes a proposition as

²For some exceptions, see Viebahn and Vitter (2016).

³Her account actually appeals to ordering sources, which are even more complicated.

an argument, and requires that proposition be true at *all* the worlds that the salient ordering ranks highest. And a modal like “may” requires its propositional argument to be true at *some* world that the salient ordering ranks highest. For example, (1) is true iff you give Bill a hundred dollars at *all* the worlds that the salient ordering ranks highest.

The rest of this paper will assume that modal terms do have an argument-place for something like an ordering. This assumption is an empirical claim about natural language semantics, not a question we can settle from the armchair. But it is highly plausible. It is much more plausible than other parts of Kratzer’s accounts of the semantics of modals. It’s even common ground among philosophers who reject many other parts of Kratzer’s semantics. For example, it’s incorporated into the dynamic approaches that of Malte Willer (2014) develops and, and into the relativist approach that Niko Kolodny and John MacFarlane (2010) develop.

I intend the word “ordering” as a theoretical term that refers to whatever the context supplies that distinguishes the different uses of (1). So the only claim I take from Kratzer is that the context supplies something that does that. I use the term “ordering” rather than the more orthodox term “ordering source” in hopes of emphasizing just how wide a range of views could accept the claim that I’ll be discussing.

1.2 Integrating metaethics

Moral realists should be open to the Kratzerian claim that modal terms have an argument place for an ordering. They shouldn’t make their view hostage to the vicissitudes of natural language semantics; they should be ready to accept what the linguists tell them is true. So throughout the rest of the paper, I’ll just assume the Kratzerian claim. If you don’t accept it, think of this paper as a hypothetical investigation of how moral realists could best accept the Kratzerian claim if it’s true.⁴

The Kratzerian claim initially appears to threaten some important debates in metaethics. That is, the debates seem much less pressing if the claim is true. Consider, for example, the debate between naturalists like Foot (2001) and Schroeder (2007) and non-naturalists like Parfit (2011) and Scanlon (2014). Their disagreement is a metaphysical one, about the nature of moral facts: whether they are natural facts. At first, it’s hard to see how a language that contains terms with Kratzer’s semantics has the expressive power to formulate this debate. Given that semantics, the fact that I have to give Bill a hundred bucks is 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. But naturalists and non-naturalists are not disagreeing about the nature of rankings on possible worlds! (Nor about the nature of sets of propositions!)

⁴For a sophisticated discussion of how the Kratzerian claim could be false, see Viebahn and Vitter (2016).

Fortunately, though, it is straightforward to express the debate about moral naturalism even even we accept the Kratzerian claim. 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 captures the demands of morality; they might agree on every normative question. 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 will 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.

Then naturalists and non-naturalists both have an account of what distinguishes the moral use of modal terms from other uses of modal terms. The moral use carries the commitment that the salient ordering is the moral standard. To illustrate, go back to (1).

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

The moral use of (1) might communicate a singular proposition about an ordering.

- (*) \mathbf{o}_1 is the moral standard and $[\forall w: w \text{ is one of the worlds that } \mathbf{o}_1 \text{ ranks highest}]$ (You give Bill a hundred dollars in w)

The second conjunct of this proposition is metaphysically *innocuous*. Even an anti-realist like Mackie can acknowledge that it’s true. It is the first conjunct that carries a metaphysically *loaded* commitment, in including the kinds of properties that Mackie takes to apply to nothing. So Mackie might say that (1) is false because the first conjunct is false, even while acknowledging that the second conjunct could be true. In other words: the metaphysically loaded conjunct is the one that contains the property that naturalists and non-naturalists are disagreeing about, the property that’s part of mind-independent moral reality. And moral skeptics are concerned to show that we don’t know the metaphysically loaded part. They should be open to the possibility that we could know the metaphysically innocuous commitment.

More generally, part of what (1) asserts is metaphysically loaded, and part of it is metaphysically innocuous – at least if moral realism is true. It’s easy to distinguish the two parts in (*), because (*) is a conjunction of a metaphysically loaded proposition with a metaphysically innocuous

propositions. But (*) is only one candidate for what the moral use of (1) would assert, if moral realists are right. Other candidates include:

- $[\exists x: x \text{ is the moral standard}] ([\forall w: w \text{ is one of the worlds that } x \text{ ranks highest}] (\text{You give Bill a hundred dollars in } w))$
- $[\text{the } x: x \text{ is the moral standard}] ([\forall w: w \text{ is one of the worlds that } x \text{ ranks highest}] (\text{You give Bill a hundred dollars in } w))$

Each of these propositions contains a metaphysically loaded part (the restrictor of the quantifier), and a metaphysically unloaded part (the material after the restrictor).

2 The metaphysically loaded part as a presupposition

I claim that the metaphysically loaded part of what (1) asserts must be interpreted as part of a *presupposition*. My argument for this claim will be important in the rest of the paper, because it has important consequences for what it takes for a knowledge report to be true.

The presuppositions of an utterance are information that that utterance (a) is somehow associated with, but (b) is somehow interpreted as being part of the background assumptions in the conversation. Examples include:

- (i) It wasn't Bill who danced.
- (ii) They don't like the president.
- (iii) Jenny maybe also danced.

Each of these sentences are somehow associated with two propositions, with one proposition interpreted as backgrounded and not-at-issue. For example, (i) is associated with the backgrounded, not-at-issue proposition that someone danced. You wouldn't use (i) unless you accept that someone danced. But the main point of using (i) is to communicate that Bill didn't dance.⁵

One of the central hallmarks of presuppositions is that they *project* from canonical embeddings. A commitment of a sentence S projects when uses of S, $\lceil \text{not-}S \rceil$, $\lceil \text{it's } 20\% \text{ likely that } S \rceil$, and $\lceil \text{if } S, \text{ then } P \rceil$ are all associated with that commitment. If I say "it was Bill who danced", I'm committed to the proposition that someone danced. And if I say "it wasn't Bill who danced", I'm still committed to the very same proposition: the proposition that someone danced.

I claim that, if moral discourse has a metaphysically loaded part, the metaphysically loaded part has to be part of what projects. Now before digging into the details, we should note one striking difference between the moral case and the examples of projection in (i)–(iii). A competent speaker-hearer who understands (i)–(iii) immediately recognizes what commitments they carry, and so can immediately see that some commitments project. And that is why (i)–(iii) are effective ways to introduce the idea

⁵Craige Roberts (2012) is the inspiration for the label 'not-at-issue'.

of commitments that project: competent speaker-hearers can immediately recognize that they project. But in other cases, it will be harder for competent speakers to immediately recognize what projects.⁶ It will be especially hard when the commitments are about theoretical posits that are not obvious to every competent speaker. And the metaphysically loaded part of what (1) asserts *is* about a non-obvious theoretical posit: the ordering that Kratzerians posit. So in trying to figure out if the metaphysically loaded part projects, we have to proceed indirectly. We have to see if there is data that can only be explained by supposing that the part projects. We shouldn't expect ordinary speakers to tell one way or another.

And there *is* decisive indirect evidence that the metaphysically loaded part does project. The evidence concerns examples like (2) and (3).

(2) You may not let the drowning child drown.

(3) You must save the drowning child.

We're ready to infer (3) from (2), and to infer (2) from (3). These two sentences illustrate what we can call the 'duality' of "may" and "must": that $\lceil \text{Must}(p) \rceil$ is true iff $\lceil \neg \text{May}(\neg p) \rceil$ is true.

One of the important virtues of Kratzer's semantics is that it crisply explains 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.

I'll show that the only way for "may" and "must" to be duals if is for the metaphysically loaded part to project. I'll introduce the argument with the hypothesis that the moral use of a modal communicates an existential proposition about the moral standard, and then generalize. Suppose for reductio that the metaphysically loaded part doesn't project. Then the moral use of (2) would communicate that

(2lf-*without-Projection*) $\neg([\exists x: x \text{ is the moral standard}]([\exists w: w \text{ is one of the worlds that } x \text{ ranks highest}] (\text{You don't let the drowning child drown in } w)))$

And the moral use of (3) would communicate that

(3lf) $[\exists x: x \text{ is the moral standard}]([\forall w: w \text{ is one of the worlds that } x \text{ ranks highest}] (\text{You save the drowning child}))$

And (2lf-*without-Projection*) does not entail (3lf). For one thing, (3lf) entails that something is the moral standard, and (2-*without-Projection*) does not.

So the only way that (3) would follow from (2) is if the metaphysically loaded part projects above the negation.

⁶See Craig Roberts (ms) for an extended discussion of one such example.

(2lf-*with-Projection*) $[\exists x: x \text{ is the moral standard}] \neg([\exists w: w \text{ is one of the worlds that } x \text{ ranks highest}] \text{ (You don't let the drowning child drown in } w))$

(3lf) does follow from (2lf-*with-Projection*). Since we think (3) follows from (2), we have to think that (2)'s metaphysically loaded part projects.

In making this argument, I've adopted one particular way for moral discourse to include a metaphysically loaded part. I've taken the moral use of sentences like (3) to communicate existential propositions about the moral standard. But there are other options, too, as we saw earlier. Maybe it communicates singular propositions, or propositions with a definite description.

But the same kind of argument will establish that *however* moral discourse includes a metaphysically loaded part, the metaphysically loaded part has to project. The crucial fact is that whatever includes the metaphysically loaded part will entail that some ordering is the moral standard. That is, it will entail that some ordering has the property that realists take to be most fundamental. Suppose, for example, that the moral use of (1) asserts a proposition that contains a definite description: something with the form "[the x : x is the moral standard] (...)". In other for that proposition, something must be the moral standard.

So we can generalize the previous argument.

MASTER ARGUMENT FOR PROJECTION

- A. If S_1 entails S_2 , then S_1 entails everything that S_2 does.
- B. In ordinary moral contexts, (2) entails (3).
- C. So in ordinary moral contexts, (2) entails everything that (3) does.
- D. The moral use of (3) entails that some ordering is the moral standard.
- E. So in ordinary moral contexts, (2) entails that some ordering is the moral standard.
- F. The metaphysically loaded part of what (2) asserts is what entails that some ordering is the moral standard.
- G. If a sentence with the form $\lceil \neg S \rceil$ entails that p , then p is part of what projects.
- I. So in ordinary moral contexts, the metaphysically loaded part of what (2) asserts is part of what projects.

I regard this argument as decisive evidence that the metaphysically loaded part of what (2) asserts is part of what projects.

The metaphysically loaded part is not just the existential proposition that something or other is the moral standard. The metaphysically loaded part figures in a substantive claim about moral reality – for example, that independent moral reality requires you to save the drowning child. Entailing that something is the moral standard is a hallmark of the metaphysically loaded part of (2). But it is not the whole of the metaphysically loaded

part of what (2) asserts. To see why, suppose that the moral use of (1) asserts the following singular proposition.

(*) \mathbf{o}_1 is the moral standard and $[\forall w: w \text{ is one of the worlds that } \mathbf{o}_1 \text{ ranks highest}]$ (You give Bill a hundred dollars in w)

The first conjunct is the metaphysically loaded part, and it entails that something is the moral standard. But the first conjunct is a much stronger claim than the mere claim that something is the moral standard. It is the claim that one particular ordering is the moral standard. But because it entails that something is the moral standard, we can use that entailment to test whether the first conjunct projects above negation. And the MASTER ARGUMENT shows that that entailment does project above negation, which shows that the metaphysically loaded part that carries that entailment must also project. In fact, it shows that the metaphysically loaded part *must* project above negation, whatever form it takes.

At this point, someone might object that there are uses of (3) that don't entail (2). For example, an error theorist might assert (3) without being committed to (2). Such a person might think they've found a reason to reject my Premise C. But they're wrong: they've just noticed an instance of perfectly general fact about presupposition. If I say "John doesn't know that it's raining", my utterance usually presupposes that it's raining. But if I first say, "it's not raining, so John doesn't..." then the utterances does not presuppose that it's raining. Presuppositions project in *ordinary* contexts, not in all contexts. The contexts where we would infer (3) from (2) are decisive evidence that the metaphysically loaded part projects.⁷

Exactly parallel arguments can be found for other constructions.

(4) It's 20% likely that you may let the drowning child drown.

(5) It's 80% likely that you must save the drowning child.

We'll often infer (5) from (4). But if the metaphysically loaded part doesn't project, (4) could be true while (5) is false. Suppose that there's a 80% chance that something is the moral standard, and a 25% chance that anything that's the moral standard will permit letting the child drown. If (4)'s metaphysically loaded part didn't project, then (4) would be true in this case. (The chances of the two propositions both being true is the product of their individual chances of being true: 25% of 80% = 20%.) But if (5)'s metaphysically loaded part doesn't project, (5) would be false. There's at most a 75% chance that anything that's the moral standard will require saving the drowning child. And 75% of 80% = 60%.

But if the metaphysically loaded parts do project, then (4) is true iff (5) is true.

⁷Someone else might object that (3) is a *conversational implicature* of (2), at least in many ordinary contexts, as Jonas Olson (2014) suggests. And if it's a conversational implicature, our tendency to infer (3) from (2) doesn't show anything special about presupposition. I can grant Olson's view, for my purposes here. Conversational implicatures can also be presuppositions; for discussion, see Robert Stalnaker (1973, 452) and Scott Soames (2009, 86–91).

(4lf) $[\exists x: x \text{ is the moral standard}]$ (It's 20% likely that $[\exists w: w \text{ is one of the worlds that } x \text{ ranks highest}]$ (You don't let the drowning child drown in w))

(5lf) $[\exists x: x \text{ is the moral standard}]$ (It's 80% likely that $[\forall w: w \text{ is one of the worlds that } x \text{ ranks highest}]$ (You save the drowning child in w))

If the metaphysically loaded part projects, there would only be one proposition embedded under each probability operator. Moreover, the embedded proposition in (4lf) is a logical contradictory of the embedded proposition in (5lf): if one is false, the other is true. So if one is 20% likely, the other has to be 80% likely. We can vindicate our antecedent convictions that we can infer (4) from (5) if the metaphysically loaded part projects. And in fact, that's the only way to vindicate our antecedent convictions.

I conclude that, if moral discourse carries a metaphysically loaded part, the metaphysically loaded part has to be part of what projects in many contexts.

My overarching claim in this section is that the metaphysically loaded part is part of a presupposition. I've just shown that it projects. Unfortunately, though, evidence that it projects is not by itself decisive evidence that it is a presupposition. There are contents that project but that are not presuppositions. Non-restrictive relatives like “who lives near me” are one example. To complete my argument, I need to show that the metaphysically loaded part has the hallmarks that distinguish presuppositions from other kinds of projective contents.

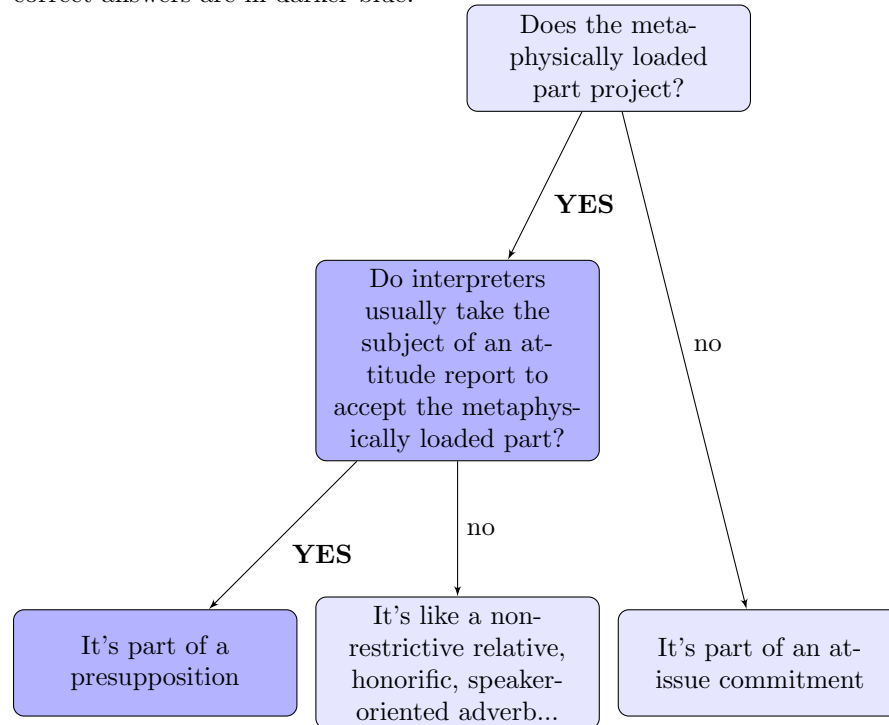
Presuppositions pattern with attitude ascriptions in a different way than other kinds of projective contents. Projective contents that are not presuppositions tend to be attributed to the *speaker* rather than the subject of the attitude verb. “Damn”, for example, is associated with a projective content – but if I say “John thinks that that damn cat is outside”, speaker-hearers often interpret *me* rather than John as accepting the projective content that's associated with “damn”. Presuppositions pattern differently. If I say “John thinks that it stopped raining”, speaker-hearers interpret John as accepting the presupposition, that it used to be raining.⁸

Attitude reports reveal that the metaphysically loaded part of moral discourse behaves like a presupposition. Imagine that John found a book that described just the ordering \mathbf{o}_1 that is the moral standard, and that he memorized the demands of the ordering. He could then have beliefs about the metaphysically innocuous part of (1); for example, the belief that $[\forall w: w \text{ is one of the worlds that } \mathbf{o}_1 \text{ ranks highest}]$ (You give Bill a hundred dollars in w). And he could have that belief without thinking of that book as having anything to do with morality – he might think that it describes the laws of some alien society. But in that case, he doesn't have a *moral* belief. He doesn't have moral belief because he doesn't accept that \mathbf{o}_1 is the moral standard. Our judgment about this case reveals that

⁸Though the data here is complicated; see Schlenker (2003), Amaral et al. (2007), and Potts (2007); see also Heim (1992) for a classic discussion.

the metaphysically loaded part does behave exactly like a presupposition: subjects of attitude reports need to accept it.⁹

This section has been filling out the following chart, showing that the correct answers are in darker blue.



At this point, I take myself to have shown that moral realists should accept a Kratzerian approach to modal terms by taking the metaphysically loaded part of moral discourse to be part of a *presupposition*.¹⁰

3 Presupposition matters for epistemology

I've just been arguing for a thesis about the semantics and pragmatics of moral discourse: that realists should take the metaphysically loaded parts of moral discourse to be *presuppositions*. The rest of this paper will show that this thesis has sweeping consequences for moral epistemology.

The presuppositional thesis has sweeping consequences for moral epistemology because presupposition triggers interact distinctively with attitude reports, including reports of what's known. To describe the distinctive behavior, let 'V' be an arbitrary attitude verb.

⁹I'm *not* claiming that it's transparent to ordinary agents that their moral beliefs involve accepting the metaphysically loaded part. I don't expect that Simon Blackburn would agree! I'm rather claiming that *if* realists are right, then his genuine moral beliefs involve accepting the metaphysically loaded part.

¹⁰I'm not the first one to suggest that modal terms carry substantial presuppositions; Stephen Finlay (2014), Dan Lopez da Sa (2008, 2015), Matthew Mandelkern (2017), Alex Silk (2016, 2017a,b) have made related suggestions.

- **FACT 1**: $\lceil A \text{ V-es that } S \rceil$ is usually true in a context only if A accepts what an utterance of S would presuppose in that context.
- **FACT 2**: $\lceil A \text{ V-es that } S \rceil$ can be true in a context even if A does **not** V what an utterance of S would presuppose in that context.

To illustrate, consider (6a).

(6a) Mary hopes that it wasn't Bill who ate the cookies.

On hearing (6a), we tend to infer that Mary accepts that someone ate the cookies – we tend to infer that she's ready to assume that someone did eat them. This fact about (6a) is an instance of the general FACT 1. Importantly, though, (6a) can be true even if Mary doesn't hope that someone ate the cookies. Maybe she's Bill's dietician, and only hopes that he's successful in avoiding temptation; she doesn't care what anyone else does. This fact about (6a) is an instance of the general FACT 2.

Given these two FACTS, the presuppositional thesis defended in §2 has sweeping consequences for moral epistemology. "Knows" is an attitude verb. So if the metaphysically loaded part is a presupposition, ordinary moral knowledge wouldn't require *knowing* the metaphysically loaded part. It would only require accepting the metaphysically loaded part. And since it's easier to *accept* a commitment than to *know* it, ordinary moral knowledge would be easier to acquire.

But before turning to the issues in moral epistemology, I need to convince you that FACT 1 and FACT 2 are both genuine facts.

- **FACT 1**: $\lceil A \text{ V-es that } S \rceil$ is usually true in a context only if A accepts what an utterance of S would presuppose in that context.

There is broad agreement that attitude reports communicate that the subject has some affirming attitude towards presuppositions of the complement. Paul Elbourne (2005), Bart Geurts (1998), Irene Heim (1992), Lauri Karttunen (1974), and Robert Stalnaker (2002) all agree on this basic point. And it's not surprising that they do agree. When you consider an arbitrary presupposition trigger embedded under an attitude report, you'll ordinarily infer that the subject of the attitude report has some affirming attitude towards the presupposition.

Stalnaker holds that attitude reports usually require the subject to *accept* the presupposition of the complement, where "to accept a proposition is to treat it as true for some reason. One ignores, at least temporarily, and perhaps in a limited context, the possibility that it is false" (Stalnaker 2002, 716). Heim, by contrast, suggests that the subject usually needs to *believe* the presupposition. Stalnaker will agree with Heim whenever the subject does believe the presupposition triggered from the complement. After all, believing a proposition is a great reason for accepting it.

But Stalnaker will disagree with Heim when the subject accepts the presupposition without believing it. Stalnaker's account is more plausible. Suppose I'm hypothetically assuming that someone stole the cookies, to figure out who would have stolen them if anybody did. I can appropriately say

“I know that it wasn’t Mary who stole the cookies” if I know that she was nowhere near the cookies. In this case, I don’t *believe* the presupposition that someone stole the cookies. So Heim’s suggestion predicts, implausibly, that my utterance is inappropriate. Stalnaker’s account, by contrast, predicts that it *is* appropriate: I accept the presupposition because I’m treating it as true for my conversational purposes. Because Stalnaker’s account is more plausible, I’ll follow him in assuming that acceptance is what matters. So I’ll assume that FACT 1 is genuinely a fact.

Now let’s turn to FACT 2.

- FACT 2: $\lceil A \text{ V-es that } S \rceil$ can be true in a context even if A does **not** V what an utterance of S would presuppose in that context.

This FACT is the fundamental reason why the §2 claim about presupposition is so important. It is the link between the issues about presupposition in §§1–2 and the questions in moral epistemology that §§4–6 will explore.

I illustrated FACT 2 with (6a).

(6a) Mary hopes that it wasn’t Bill who ate the cookies.

I noted that (6a) can be true even if Mary doesn’t hope that someone ate the cookies. FACT 2 is a generalization from this kind of case; in generalizing, I am taking (6a) to illustrate a general fact about presupposition and attitude verbs, rather than something specific about the verb *hopes* or the particular presupposition trigger in (6a).

And it is highly plausible that (6a) does illustrate a general fact about presupposition and attitude verbs. We see just the same behavior with a wide range of different kinds of attitude verbs.

(6b) John **conjectures** that it wasn’t Bill who ate the cookies.

(6c) John **is glad** that it wasn’t Bill who ate the cookies.

(6d) John **knows** that it wasn’t Bill who ate the cookies.

(6b) can be true even if John doesn’t conjecture that someone ate the cookies; (6c) can be true even if he isn’t glad that someone ate the cookies; and (6d) can be true even if he doesn’t know that someone ate the cookies. If you’re skeptical about my claim about (6d), imagine that John is assuming that the cookies were eaten to figure out who would have eaten them – though he’s in fact unsure if they were eaten. If he knows that Bill went nowhere near the cookies, (6d) would be true.

It will be important throughout the rest of the paper that $\lceil A \text{ knows that } S \rceil$ can be true even if A doesn’t know S’s presuppositions. And I’ve already given one piece of evidence that it can. But because this claim is *so* important, we should look for other pieces of evidence for it as well. Consider this discourse.

(7a) I don’t know whether someone ate the cookies. But I’ll accept for now that someone did. I definitely know that it wasn’t Bill who ate the cookies. I’ve been watching him all day, and he went nowhere near them.

(7a) is appropriate, even though I acknowledge that I don't know that presupposition. So knowing that it wasn't Bill who ate the cookies doesn't in general require knowing that someone ate the cookies.

It's worth double-checking this data point. We should make sure that the felicity of (7a) is probative about presupposition, rather than probative about this particular linguistic environment. Consider (7b).

(7b) I don't know whether Bill ate the cookies. But I'll accept for now that he didn't. * I definitely know that it wasn't Bill who ate the cookies. The plate is empty, so someone ate the cookies.

(7b) *is* defective. The third sentence asserts that I know that Bill didn't eat the cookies, and the first sentence asserts that I don't know whether he did. So the pattern in (7a) does reveal something about presupposition, rather than something about the particular linguistic environment. This difference is a striking confirmation of FACT 2.

I've been using a cleft ("it wasn't Bill who") as an example of a presupposition trigger. But other presupposition triggers pattern similarly.

(7c) I don't know whether the US still has a president and not an emperor. But I'll accept for now that it does. I definitely know that the president embarrasses me.

(7d) I don't know whether Beth danced. But I'll accept for now that she did. I definitely know that Jane also danced; I danced with her.

We should conclude that $\lceil A \text{ knows that } S \rceil$ can be true even if A doesn't know S's presuppositions.¹¹

¹¹Now there is significant heterogeneity among the class of presupposition triggers; for example, "anaphoric" triggers like *too* and *also* behave differently than non-anaphoric triggers like *stop* and *start*. Tonhauser et al. (2013) describe the differences between different presupposition triggers in illuminating detail. Interestingly and importantly, though, the class of presupposition triggers *is* unified in conforming to FACT 2.

Importantly, knowing that S does *not* require conditional knowledge, knowledge that if the not-at-issue proposition is true, the at-issue proposition is too. In some cases, it might seem to require that knowledge. For example, knowing that it was Bill who ate the cookies seems to require knowing that if someone ate the cookies, Bill ate them. However, this requirement doesn't hold generally. Knowing that it wasn't Bill who ate the cookies doesn't require knowing that if someone ate the cookies, it wasn't Bill. Just imagine that there is a gremlin who forces Bill to eat the cookies if anyone tries to eat them. The gremlin is inactive, because in fact no one tried to eat the cookies. But I don't know any of this, and I accept that someone did eat the cookies. In that case, I don't know the conditional. But because I've been watching Bill all day and accept that someone ate the cookies, I can truly say that I know that it wasn't Bill who ate the cookies. Sentences like "it was Bill who ate the cookies" pattern differently because the at-issue proposition *entails* the not-at-issue proposition: that Bill ate the cookies entails that someone ate the cookies.

There is also a more theoretical argument for this conclusion. We've noted that FACT 2 holds very generally.

- FACT 2: $\lceil A \text{ V-es that } S \rceil$ can be true in a context even if A does **not** V what an utterance of S would presuppose in that context.

And all the leading accounts of presupposition triggers in attitude reports give a highly general explanation of FACT 2. They explain FACT 2 as following from very general facts about presupposition, rather than idiosyncratic facts about particular attitude verbs like “hope” or “conjecture”.¹² So we would expect the pattern that we see with “hope” to hold generally. In particular, we would expect it to hold for verbs like “know”. In other words: the only descriptively adequate accounts will entail FACT 2, thereby entailing the crucial claim about knowledge reports, since the claim about knowledge reports is just a substitution instance of FACT 2. In fact, I see the generality of FACT 2 as being more probative for my claim about knowledge than our judgments about particular cases. Once you recognize how general the correct explanation of FACT 2 has to be, you have no choice but to accept that the same explanation will hold for “know”.¹³

4 A concrete picture of how the presuppositional thesis could be true

I take myself to know that killing is usually wrong. Moral skeptics characteristically disagree. In disagreeing, they are implicitly defending both a *conceptual* claim about what moral knowledge requires and a *substantive* claim about my epistemic position.

The presuppositional thesis defended in §§1–2 has sweeping consequences for what it takes to have moral knowledge. So it has sweeping consequences for the conceptual claims that moral skeptics rely on. In highlighting these sweeping consequences, I'll use linguistic claims about sentences like (8) as evidence for conceptual questions about what moral knowledge requires.

(8) I know that killing is usually wrong.

Why is it legitimate to use our *linguistic* judgments about *sentences* like (8) as evidence about the *non-linguistic* question of whether the skeptic is right? It's legitimate because our judgments about knowledge involve

¹²See Irene Heim (1992) for the benchmark account, and see also Márta Abrusán (2011), Luka Crnic (2011), Kai von Fintel (1999, 2008), Aynat Rubinstein (2017), Philippe Schlenker (2009), and Elizabeth Villalta (2009).

¹³Now there might be one exception to FACT 2. It would not hold for verbs like “believe”, if Heim is right that attitude reports require the subject to *believe* presuppositions of the complement. But this observation does not threaten my claim that the pattern holds generally. Belief plays a central role in the leading explanation of the general pattern. For example, Heim's explanation crucially involves the claim that the subject is required to believe presuppositions of the complement. So of course believing that S will require believing S's presuppositions – but only because V-ing that S in general requires believing that S.

dispositions to accept or reject sentences ascribing knowledge. We can make progress on understanding what moral knowledge requires by gaining a better understanding of the linguistic dispositions go hand-in-glove with our judgments about knowledge.¹⁴

This section will give one account of how we can know about the metaphysically innocuous part of moral discourse. This account isn't the only possible one. And it is an addition to the presuppositional thesis defended in the previous section. But I will use it in later sections to illustrate the sweeping upshots of the presuppositional conjecture.

The metaphysically innocuous part of moral discourse is about the demands of an *ordering*, which is a complicated object, appropriately modeled with very complicated set-theoretic constructions. So orderings don't look like the sort of things that ordinary, limited agents like you or me could now about. But consider a comparison. The *language* that I speak is also a complicated object, appropriately modeled with very complicated set-theoretic constructions. But we are willing to take me to know a great deal about my language. For example, we're willing to take me to know that "salad" means salad in my language. And that piece of knowledge is implicitly knowledge about my particular language – I can still acknowledge that "salad" didn't mean salad in seventeenth-century English.¹⁵

I will use knowledge of my own idiolect as my model for knowledge of orderings. I have a set of dispositions to apply the word "salad" to some

¹⁴Though I focus on linguistic judgments, I am not convinced that linguistic judgments have a privileged role in providing evidence for my presuppositional hypothesis. For one thing, the distinction between what's at-issue and what's not-at-issue may primarily be a distinction at the level of thought rather than language. That is, when reasoning by myself about some question, I may implicitly treat some propositions as at-issue and others as not-at-issue, independently of the language that I happen to speak. (Simons et al. (2017) are at the cutting edge of developing an approach that could be interpreted in that way.) In that case, the evidence from §1 would be direct evidence about the concept of moral knowledge that we employ, rather than indirect evidence from the language that we speak about the concept of moral knowledge that we deploy.

Even if the §1 arguments only give us indirect evidence, evidence about the language we speak, they would still give us evidence that moral skeptics are not using the concept of moral knowledge that we actually employ. There is still room for lively debate about the upshot of this conclusion. It's very much like the debate over linguistic evidence for epistemic contextualism, with figures like Hilary Kornblith (2000) arguing that the evidence has limited relevance, and Keith DeRose (2005) responding.

¹⁵Importantly, though, the sort of knowledge that I'm describing is not essential for semantic competence with your idiolect. For example, maybe you can speak a idiolect without being in a position to think about semantic concepts like *means-in*. And some philosophers are very interested in explaining our linguistic knowledge without us thinking about those sort of semantic concepts. (Davis Lewis is one example, in his "Languages and Language" (1975).) I am interested in a different question. I am focusing on people who already have the semantic concept *means-in*. I want to explore how they have special and privileged knowledge of what falls under that concept, in their own idiolect.

objects and not to others. Present me with an object, and I can usually tell you if it's in the extension of "salad" in my idiolect. There is some important relationship between this set of dispositions and my semantic knowledge, my knowledge of what "salad" means. One possibility is that my set of dispositions *grounds* much of my semantic knowledge. Given this possibility, "salad" means what it does in my idiolect because of my dispositions. (In saying that my dispositions ground the meanings of some terms, I'm not saying that they ground the meaning of all terms – maybe they don't ground what "water" means, for example.) Another possibility is that the set of dispositions is just a hallmark of this semantic knowledge. Given this possibility, my dispositions just reveal when I do have semantic knowledge.

David Lewis (1970) popularized the technique of *Ramsifying*. The key idea is to take individuals to be thinking about whichever idiolect best fits their linguistic dispositions. If I'm disposed to apply 'red' to objects whose wavelength is between 615nm and 695nm, then the meaning of 'red' in my idiolect is (roughly) the property of falling within that range.

Then my linguistic dispositions can sometimes ground knowledge of what different terms mean in my idiolect. Suppose that my dispositions with some term wholly determine what it means in my language. Then my dispositions with it are guaranteed to deliver true beliefs about its meaning. And this guarantee also explains why my true beliefs about its meaning will amount to knowledge. Now in many, perhaps almost all cases, my dispositions don't wholly determine what a term means – external factors may play a significant role. But in the cases when my linguistic dispositions play a significant enough role, those dispositions can ground knowledge.¹⁶

Let's now turn back to the moral case, using the Ramsifying suggestion about idiolects as a model. I'm disposed to make a range of moral judgments: to think that killing is usually wrong, but that killing in self-defense is usually permitted, though not if I saw a way to defuse the threat without killing, and on and on. There is some ordering that best fits those judgments. I've just suggested that linguistic dispositions can ground an ability to think about a particular idiolect. I propose that my *moral* dispositions can similarly ground an ability to think about a particular ordering.

I further propose that my moral dispositions also allow me to *know* about the demands of the ordering. I suggested earlier that my linguistic dispositions can allow me to know about a particular idiolect in those cases where they're guaranteed to be a reliable guide to it. Similarly, my moral dispositions can allow me to know about a particular ordering in those cases where they're guaranteed to be a reliable guide to it. I'm here waving my hands to indicate a concrete picture about how I could know about the demands of a particular ordering. I hope that the analogy with linguistic dispositions has put enough flesh on the concrete picture for you to understand what sort of picture I'm waving at.

This concrete picture should be sharply distinguished from the presuppositional thesis defended in §2. The presuppositional thesis could be true

¹⁶See Lewis (1983, 1984).

even while the concrete picture is false. As should already be apparent, the concrete picture embodies some substantive claims that you might reject. And you can reject those substantive claims without rejecting the presuppositional thesis. My goal in introducing this concrete picture is to illustrate one way that the presuppositional conjecture could be true, so that we can appreciate some of its philosophical payoffs. Even if you regard the concrete picture as misguided, you can still rely on it to illustrate the difference that the presuppositional thesis makes in moral epistemology.

5 Knowledge of orderings can suffice for moral knowledge

I've just finished sketching a concrete picture where I can know about the demands of a particular ordering. I'll use this concrete picture to illustrate the philosophical payoffs of the presuppositional thesis that §§1–2 defended. It turns out to be surprisingly easy to have some kinds of moral knowledge.

§1 noted several different ways for moral realists to accept a broadly Kratzerian semantics for modals. One option is to take moral discourse to communicate singular propositions about particular orderings, like \mathbf{o}_1 . In showing that the metaphysically loaded part has to be part of a presupposition, I would then be showing that the proposition that \mathbf{o}_1 is the moral standard is a presupposition of moral discourse. Since knowledge reports don't require knowledge of the presupposition, the moral use of "I know that killing is usually wrong" would ordinarily communicate that:

SINGULAR GLOSS I know that \mathbf{o}_1 usually forbids killing and accept that \mathbf{o}_1 is the moral standard, and \mathbf{o}_1 is the moral standard.¹⁷

My concrete picture is that my dispositional moral judgments can ground knowledge of particular orderings, like \mathbf{o}_1 . So the concrete picture explains how I can satisfy the first requirement in the SINGULAR GLOSS.

But if my knowledge of \mathbf{o}_1 is grounded in my dispositional moral judgments, I also accept \mathbf{o}_1 to be the moral standard. Remember that the property *being-the-moral-standard* is the property that realists posit as what distinguishes moral beliefs from other kinds of modal beliefs, like beliefs about what prudence demands. So the ordering that best fits my moral beliefs has to be an ordering that I implicitly accept to be the moral standard. So I'll also satisfy the second requirement of the SINGULAR GLOSS. So given the concrete picture from §4, I can have moral knowledge if I have true dispositional moral beliefs. That is, having the relevant range of true beliefs *can itself be sufficient* for moral knowledge. Moral skeptics are assuming

¹⁷I'm here abstracting further from some aspects of Kratzer's semantics. Let "forbid" abbreviate "doesn't happen in any of the worlds that \mathbf{o}_1 ranks highest.

And this knowledge report requires that it's true that \mathbf{o}_1 is the moral standard, even though it doesn't require *knowing* that it's true. Presuppositions project through "knows": in order for $\ulcorner A$ knows that $S \urcorner$ to be true, the presuppositions of S have to be true; Lauri Karttunen (1974) has the classic discussion.

otherwise. So they are assuming some conceptual claims about what moral knowledge would take that are a poor fit for the presuppositional thesis.

Now the SINGULAR GLOSS is just one way that the presuppositional thesis might be true. The other ways for it to be true are ways where our moral attitudes are about quantificational propositions about the moral standard. For example, the moral use of “I know that killing is usually wrong” might communicate that:

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

These quantificational glosses will have the same consequences for moral epistemology that the SINGULAR GLOSS does. Since the metaphysically loaded part (designated by ‘ x is the moral standard’¹⁸) is part of a presupposition, moral knowledge doesn’t require knowing about it.

Now it is harder to see how knowledge about a particular ordering could suffice for knowledge of the EXISTENTIAL GLOSS. To see how it could, let’s begin with some very general observations. If I believe that a is F and believe that a is G , I’ll also believe that an F is G , as long as I realize that both beliefs about a are about the same object.¹⁹ A similar point holds for knowledge. If I know that a is F and know that a is G , I can also know that an F is G , as long as I realize that both pieces of knowledge are about the same object. For example, if I knew that that \mathbf{o}_1 is the moral standard and know that \mathbf{o}_1 usually forbids killing, then I can know that $[\exists x: x \text{ is the moral standard}]$ (x usually forbids killing).

Importantly, though, it’s even easier to know that $[\exists x: x \text{ is the moral standard}]$ (x usually forbids killing) if the metaphysically loaded part is a presupposition. I’d be in a position to know that proposition if I merely *accept* that \mathbf{o}_1 is the moral standard while knowing that \mathbf{o}_1 usually forbids killing – because knowledge reports don’t require knowledge of the presupposition. Just like knowledge of ordering is sufficient for the SINGULAR GLOSS being true, knowledge of orderings is also sufficient for the EXISTENTIAL GLOSS being true.

To appreciate this crucial point, consider the benchmark account of presupposition triggers under attitude reports that Irene Heim (1992) developed. For her, attitude reports implicitly restrict attention to worlds where the presupposition is true. To evaluate if “I know that killing is usually wrong” is true, we restrict attention to worlds where the ordering I accept to be the moral standard (\mathbf{o}_1 , say) is the moral standard. Then we ask if I know throughout those worlds that \mathbf{o}_1 usually forbids killing. If I do, then the knowledge report is true. More generally, given the presuppositional thesis and Heim’s account, knowing that \mathbf{o}_1 usually forbids killing

¹⁸Take ‘ x is the moral standard’ to designate a propositional function, from objects to the proposition that that object is the moral standard. It’s not a genuine proposition – that’s why I formulate the whole paper in terms of the metaphysically loaded *part* of moral discourse, rather than a metaphysically loaded *proposition* associated with moral discourse.

¹⁹That is, I’ll believe that an F is G when the beliefs about x need to be “coordinated”, to drop on the concept that Kit Fine (2009) extensively discusses.

while accepting that \mathbf{o}_1 is the moral standard will be sufficient for knowing any quantificational proposition about the moral standard usually forbidding killing. So whatever form the presuppositional thesis takes, knowledge of orderings will be sufficient for moral knowledge.

6 How moral skepticism could be a pseudo-problem

§5 introduced a possible picture of what moral knowledge requires. This section will briefly indicate some philosophically important upshots of this picture. Once I've done that, I'll return to the relationship between the §5 picture and the presuppositional thesis from §§1–2. As we'll see later, it might be possible for the presuppositional thesis to be true without leading to the §5 picture. I myself think that this possibility isn't genuine – I think that the presuppositional thesis *does* lead to the §5 picture. But before exploring the relationship between the two in full detail, I want to convince you that the §5 picture is philosophically interesting.

Let's start with a moral disagreement, like the disagreement between Jefferson Davis and Frederick Douglass about the permissibility of slavery. Non-skeptical realists will think that one of them knows what they believe. That is, they'll take (1a) to be true.

(1a) Either Davis knows that slavery is morally permissible, or
Douglass knows that it's not morally permissible.

And moral skeptics deny that (1a) is true. Or, more carefully, skeptics about our knowledge about slavery will deny that it's true. That's just what it is to be a skeptic about some moral topic – denying that anyone has (substantive) moral knowledge about it.

As noted earlier, there are a wide range of problems that motivate moral skepticism. Maybe pervasive moral disagreement dooms moral knowledge about some topics; maybe empiricist strictures on knowledge do knowledge; maybe the dispensability of moral explanation do; and so on.²⁰

These problems all purport to show that, whatever propositional attitudes we have to the metaphysically loaded part of moral discourse, those propositional attitudes cannot amount to knowledge. But given the §5 picture, moral knowledge doesn't require *knowledge* about the metaphysically loaded part of moral discourse. It only requires *accepting* something about the metaphysically loaded part, while *knowing* about the metaphysically innocuous part. The arguments that moral skeptics develop simply do not threaten knowledge of the metaphysically innocuous part of moral discourse. So since the only kind of knowledge required for moral knowledge is knowledge of the metaphysically innocuous part, the skeptics' arguments simply do not threaten moral knowledge. They're posing pseudo-

²⁰David Hume (1739) leveraged empiricist strictures in arguing that moral rationalism leads to skepticism. Sextus Empiricus (2000) argued from disagreement; Katia Vavova (2014) and Kieran Setiya (2012) both have illuminating discussions of arguments from disagreement. Gilbert Harman (1977) appeals to the dispensability of moral facts in scientific explanation.

problems, in the sense that skeptics are mistaken about what it takes to have moral knowledge, rather than being mistaken about whether we meet some agreed-upon criteria for having moral knowledge.

Here’s a concrete illustration of this point. Given the SINGULAR GLOSS on moral knowledge, (1a) is true iff

(+) Either Davis knows that \mathbf{o}_1 permits slavery while accepting that \mathbf{o}_1 is the moral standard, and \mathbf{o}_1 is in fact the moral standard, or Douglass knows that \mathbf{o}_1 forbids slavery while accepting that \mathbf{o}_1 is the moral standard, and \mathbf{o}_1 is in fact the moral standard.

Importantly, neither disjunct requires *knowing* anything about the metaphysically loaded part of moral discourse. It only requires *accepting* something about the metaphysically loaded part. Since moral skeptics are trying to show that we don’t *know* about the metaphysically loaded part, their arguments just can’t show that (+) is false.

Moreover, in fact (+) is plausibly true. Given the concrete picture from §4, what matters for moral knowledge is having true dispositional moral beliefs. So (+) will be true if either Davis or Douglass has true dispositional moral beliefs about slavery. And it’s plausible that one of them will. In fact, it’s possible to sharpen the point. Take all the possible combinations of dispositional moral beliefs about slavery. Davis and Douglass had different dispositional beliefs about slavery than Aristotle did, who in turn had different dispositional moral beliefs than Genghis Khan. Then consider the long disjunction that puts together agents with each of the possible dispositional moral beliefs about slavery.

(1b) Either Davis knows ..., or Douglass knows ..., or Aristotle knows..., or Genghis Khan knows ..., or ...

Since this disjunction includes people with all possible combinations of dispositional moral beliefs, we know that one of the people will have true dispositional moral beliefs about slavery. And if the presuppositional thesis is true, those true dispositional moral beliefs suffice for knowledge about slavery. So one of the disjuncts in (1b) is guaranteed to be true. Skeptics about moral knowledge about slavery have to deny that (1b) is true. So if the presuppositional thesis is true, those skeptics have to be posing pseudo-problems.²¹

²¹This paper has focused on modal terms like “permissible” and “must”, and ignoring terms like “good” and “reason”. So I haven’t discussed skepticism about our knowledge of what’s good, or our knowledge of our reasons. But it would be surprising if “good” or “reason” patterned very differently. There seem to be conceptual connections between modal terms and the other two kinds of terms. If “you morally ought to call” is true, “calling is morally good” also seems to be true. And if “you morally ought to call” is true, “there is moral reason to call” also seems to be true. (Aaron Sloman (1970) and Aynat Rubinstein (2014) discuss related points.) There isn’t space here to argue for a parallel presuppositional account of “good” and “reason” – but in light of the connections between modal

It is helpful to contrast this answer to moral skeptics with a superficially similar answer to external world skeptics . Consider a disjunction attributing knowledge about the external world:

(1c) Either I know that the current president of the US is bald,
or my neighbor knows that the current president isn't bald.

Since knowledge attributions don't require knowledge of the presupposition (1c) can be true even if I don't know the presupposition that there is a current president. (1c) only requires accepting that presupposition. And one of the two disjuncts is true, since they are logical contradictories. It thus seems like (1c) is true. But if (1c) is true, skeptics about the external world are incorrect.

Despite initial appearances, this answer is not the same as my answer to the moral skeptic. Even though (1c) doesn't require knowledge of the presupposition, it does require knowledge of the foregrounded, at-issue commitment. If the foregrounded, at-issue commitment isn't knowledge about the existence of a president, it is about a particular individual: Donald Trump. And if it's about him, (1c) is true only if I know that he's bald, or if my neighbor knows that he's not bald. So its truth does require that we know about an object in the external world, which is just the knowledge that the external world skeptic will contest. By contrast, (1a) and (1b) require knowledge of something that's metaphysically innocuous: something like an ordering. Moral skeptics are not concerned to contest that knowledge. They are rather concerned to show that we do not know about the metaphysically loaded part of moral discourse.

So the key difference between (1a-b) and (1c) is about what is foregrounded and at-issue. In (1c), the at-issue proposition associated with the complements is a proposition about the external world: either an existential proposition about presidents, or a proposition about Trump himself. The external world skeptic contests knowledge of both propositions. Now one way to answer the external world skeptic would be to adopt a very radical hypothesis: that the at-issue propositions are propositions about *logical constructions out of sense data*, which are presupposed to correspond with the external world. That radical hypothesis would combine with my ideas about presupposition to answer the external world skeptic. Now that radical hypothesis is, I think, implausible as an empirical hypothesis about natural language. However, that radical hypothesis has the same structure as the idea about moral discourse that I've defended in this paper. (In both cases, the at-issue proposition is much easier to know than the not-at-issue proposition.) I think the presuppositional hypothesis about moral discourse is highly plausible; in fact, I've been concerned to argue that there is decisive evidence for it.

terms and "good" and "reason", it would be quite surprising if a presuppositional account were correct for modal terms but not for "good" or for "reason".

7 Does the presuppositional thesis really have philosophically interesting consequences?

I've suggested that the presuppositional thesis from §§1–2 leads to the philosophically interesting picture of moral knowledge discussed in §§5–6. But in making this suggestion, I've been glossing over one possibility where the presuppositional thesis *wouldn't* lead to that picture.

The presuppositional thesis wouldn't lead to the §5 picture in contexts where the metaphysically loaded part of moral discourse is *locally accommodated* in knowledge attributions. To introduce those contexts, think back to sentences like “I hope that it wasn't Bill who ate the cookies”. As noted earlier, it can be appropriate to use that sentence without hoping that someone ate the cookies. But in some contexts, using that sentence *would* communicate a hope that someone ate the cookies and that Bill didn't. For example, you might use it if you hope that the cookies aren't around to tempt Bill anymore but also hope that he didn't eat them. In this case, we would say that the presupposition is *locally accommodated* under “hopes”, because the hope is partially directed at the presupposition. Such cases are exceptional, but they are possible.

For entirely general reasons, then, we would expect to find contexts where the metaphysically loaded part of moral discourse is locally accommodated in knowledge reports. And in such contexts, the §5 picture of moral knowledge would be mistaken: moral knowledge would require knowing the metaphysically loaded part, and not merely accepting it. As a result, the philosophically interesting consequences of the presuppositional thesis follow only in the contexts where the metaphysically loaded part is *not* locally accommodated. And since propositions that are at-issue in the conversation are not locally accommodated, the interesting consequences don't follow when the metaphysically loaded part is at-issue. So even if the presuppositional thesis is correct, moral skeptics can argue that the metaphysically loaded part of moral discourse is at-issue and locally accommodated in knowledge attributions.

I introduced this paper with the claim that moral epistemologists have been focusing on too narrow a range of questions. We're now in a position to appreciate the full range of questions that matter. If the presuppositional thesis is true, it matters enormously whether the metaphysically loaded part is at-issue. In order for moral skeptics to press their challenge, they have to show that it is. So a new and unappreciated question is:

NEW QUESTION: Can moral skeptics show that the metaphysically loaded part is at-issue?

Given the presuppositional thesis, moral skeptics have to address the NEW QUESTION before addressing the TRADITIONAL QUESTION.

TRADITIONAL QUESTION: Can moral skeptics show that we're ignorant about the metaphysically loaded part?

If the answer to the NEW QUESTION is ‘no’, then the answer to the TRADITIONAL QUESTION is irrelevant, because moral knowledge doesn't require

knowing the metaphysically loaded part.

At this point, you might think that the answer to the NEW QUESTION is ‘yes’, and indeed that it’s easy to see that the answer is ‘yes’. The skeptic intends to be talking about the metaphysically loaded part of moral discourse. So the metaphysically loaded part becomes at-issue when speaker-hearers recognize her intention.

But this easy answer to the NEW QUESTION is philosophically uninteresting. Being at-issue is conversation-relative: something can be at-issue in one conversation but not-at-issue in another. Acceptance is also conversation-relative; something can be accepted in one conversation but not another. (To accept a proposition is to treat it as true for some reason. And the goals of the conversation determine your reasons for treating a proposition as true, and different conversations have different goals.) Moreover, there are conceptual ties between being at-issue and being accepted. A conversation where you accept a proposition is a conversation where you’re disposed to treat the proposition as not-at-issue. In accepting it, you’re treating it as settled fact for the moment. And a conversation where you’re disposed to treat a proposition as at-issue is a conversation where you’re *not* disposed accept it. In treating it as at-issue, you’re not treating it as settled fact.

The easy answer to the NEW QUESTION doesn’t illuminate the key philosophical questions, because it appeals to a reason against accepting a proposition that tell us nothing about the subject-matter. Take the proposition that there is no largest prime. Though I believe that proposition, there are contexts where there are compelling reasons against accepting it – for example, contexts where I’m trying to cooperatively reason with someone who thinks there *is* a largest prime. Moreover, since $\lceil \neg p \rceil$ and $\lceil I \text{ know } p \rceil$ can’t both be true, the envisaged context is one where I shouldn’t accept that I know that there is no largest prime– even though, of course, I really do know.

Being cooperative is thus a philosophically uninteresting reason to treat a proposition as at-issue and *eo ipso* not accept it. It doesn’t tell you anything about the domain under discussion. Crucially, though, the impression that the NEW QUESTION has an easy answer involves just this sort of cooperativeness – cooperative deference to the skeptics’ intentions. So we should set the easy answer aside, as not revealing anything special about morality. We should instead ask if there are other reasons (besides cooperativeness) to treat the metaphysically loaded part as at-issue. And asking that question is asking for reasons against accepting the metaphysically loaded part of moral discourse. I’ll say that we’re looking for *domain-specific* reasons against accepting the metaphysically loaded part and for treating it as at-issue. We’re setting aside what I’ll call *state-given* reasons, like those grounded in a state of cooperative engagement with someone.

An answer to the NEW QUESTION thus rests on the correct theory of the domain-specific reasons that bear on accepting the metaphysically loaded part of moral discourse. If those reasons on balance favor accepting it, they also favor treating it as not-at-issue; the answer to the NEW QUESTION would then be ‘no’, and moral epistemology would look very different than

philosophers traditionally assume. I'll briefly indicate two theories about the relevant domain-specific reasons, to illustrate how the NEW QUESTION is very different from the TRADITIONAL one.

Both theories will claim that the only common norm on acceptance is coherentist: that the only kind of criticism that can always be levered at someone's states of acceptance is that the states are together incoherent, either logically or probabilistically. Accepting a proposition is treating it as true for some contextually salient goal. And you are guaranteed to be in a worse place to achieve that goal if what you accept is incoherent, for familiar Dutch-book reasons. Coherence is thus a plausible norm on acceptance, whatever your goals. Both theories will also allow that some goals impose additional norms on acceptance. That is, they allow that we can appropriately criticize even a coherent acceptance state if that state is a poor way to achieving the salient goal. But those criticisms are appropriate only because of the goal; there isn't any other common norm on acceptance besides coherence. There isn't any other common norm because of the range of goals we might have in accepting a proposition. For example, you can accept a proposition that you don't believe to reason with someone who does believe it – and coherence is the only plausible norm on what you're doing when you do that.

The two theories of acceptance differ in their accounts of our goals in moral contexts. The first theory claims that our goals in moral contexts do *not* impose any further norms beyond coherence. Though this further claim is very substantive, there are already philosophers who accept it, including Norman Daniels (1979), John Rawls (1975), and T. M. Scanlon (1998, 2002). These philosophers accept this substantive claim by reflecting on first-order normative inquiry; they think it best fits the way we actually reason.

This first theory supports a distinctive answer to the NEW QUESTION.

NEW QUESTION: Can moral skeptics find domain-specific reasons for treating the metaphysically loaded part as at-issue?

Domain-specific reasons for treating the metaphysically loaded part as at-issue are domain-specific reasons against accepting it. Given the first theory, the skeptic can only provide such reasons by showing that her target's acceptance states are incoherent, either logically or probabilistically. Though many people's acceptance states are logically or probabilistically incoherent, it's possible for someone to have logically and probabilistically coherent acceptance states. For such a person, the answer to the NEW QUESTION is 'no'. And if the answer to the NEW QUESTION is 'no', the answer to TRADITIONAL QUESTION is irrelevant; arguments for moral skepticism would then be posing pseudo-problems.

And, crucially, a coherentist answer the NEW QUESTION has very different liabilities than a coherentist answer to the TRADITIONAL QUESTION.

TRADITIONAL QUESTION: Can moral skeptics show that we're ignorant about the metaphysically loaded part?

A coherentist answer to the TRADITIONAL QUESTION would be that coherence is sufficient for *knowledge* of the metaphysically loaded part. A defense of this answer would either involve defending coherentism about knowledge in general – which is hard to do²² – or articulating special features of the moral case that make coherence sufficient for moral knowledge – which is also hard to do.²³ So it is *hard* to give a coherentist answer to the TRADITIONAL QUESTION. By contrast, it is comparably *easy* to give a coherentist answer to the NEW QUESTION; it doesn't require showing that coherence is sufficient for knowledge. It only requires showing that our goals in moral contexts impose no further constraints beyond coherence.

I now turn to the second theory, which puts a heavier emphasis on the role that moral judgments play in our social life together. It focuses in particular on the reactive attitudes: resentment, indignation, and the like. Its distinctive claim is that the reactive attitudes give us sufficient domain-specific reasons for accepting the metaphysically loaded part of moral discourse. This distinctive claim rests on two key theses. The first key thesis is that the rationality of the reactive attitudes depends on our moral beliefs: that I can rationally resent your stepping on my foot only if I believe that your stepping on my foot was morally wrong.²⁴ And the second key thesis is that the reactive attitudes are an ineliminable part of ordinary human life; giving them up would require abandoning what Strawson (1974) calls the participant stance. So the second theory predicts that our reactive attitudes combine with coherence constraints to provide sufficient domain-specific reasons for accepting the metaphysically loaded part. For example: they provide domain-specific reasons for accepting that independent moral reality forbids stomping on my foot.

If the second theory is right, moral skeptics are also posing pseudo-problems, at least when we're occupying the participant stance. Domain-specific reasons for taking the metaphysically-loaded part to be at-issue would be reasons for not accepting it. But creatures like us who occupy the participant stance always have sufficient domain-specific reasons for accepting it. We can only have domain-specific reasons for treating it as at-issue when we step outside the participant stance. But those reasons disappear when we step back into the participant stance. Since we cannot stably abandon the participant standpoint, we cannot stably deny that we have moral knowledge.

Crucially, though, this Strawsonian answer is much more compelling as an answer to the NEW QUESTION than as an answer to the TRADITIONAL QUESTION. As an answer to the TRADITIONAL QUESTION, it would require arguing that psychological indispensability is sufficient for knowledge: that the fact that I'll continue to accept that stomping on my foot is morally wrong grounds knowledge that it's wrong. And psychological indispens-

²²Though see Brink (1989).

²³Though see Scanlon (2014).

²⁴Joseph Butler (1729) and John Rawls (1971) both hold that the reactive attitudes are only *psychologically possible* for agents with the corresponding beliefs; R. Jay Wallace (1998) argues for the weaker and more plausible claim about rationality.

ability doesn't in general ground knowledge. (For example, we might be unable to stop accepting that some events are absolutely simultaneous, independent of reference frame – but we still don't *know* that they are!) But psychological indispensability might well be sufficient for acceptance, since the norms on acceptance are less demanding than the norms on knowledge.

These two theories illustrate how moral epistemology would be very different if the presuppositional thesis is correct. Given each of them, the NEW QUESTION raises very different issues than the TRADITIONAL QUESTION. Unfortunately, there isn't space here to survey all the possible theories of the domain-specific reasons for and against accepting the metaphysically loaded part. But we know that the NEW QUESTION must raise different issues than the TRADITIONAL QUESTION does. The TRADITIONAL QUESTION is about knowledge, while the NEW QUESTION is about acceptance – and acceptance is a different and less demanding state.

8 A case study in disagreement

It's helpful to close with a more concrete philosophical dispute. Consider an argument from disagreement, in Katia Vavova's formulation:

1. The correct response to peer disagreement is agnosticism.
 2. There is a lot of peer disagreement about morality.
 3. Therefore, we should be agnostic about a lot of morality.
- (Vavova 2014, 304)

Given the presuppositional conjecture, we evaluate this argument in two steps. We first ask if peer disagreement shows that the answer to the NEW QUESTION is 'yes'. If it doesn't, we ask if peer disagreement shows that we're ignorant of the *not-at-issue* part of moral discourse.

In order for the answer to the NEW QUESTION to be 'yes', peer disagreement would have to be a domain-specific reason against accepting the metaphysically loaded part of moral discourse. So we can illustrate the two theories sketched in the previous section by describing what they would say about peer disagreement.

Given the first theory, peer disagreement won't provide domain-specific reasons against accepting the metaphysically loaded part. On that theory, those domain-specific reasons come from coherence or incoherence *within what the agent herself accepts*. Since what other people accept won't lead the agent to abandon what she herself accepts, peer disagreement doesn't provide any domain-specific reasons against accepting the metaphysically loaded part, and thus no domain-specific reasons for treating it as at-issue.

Let's unpack this point more slowly, noting that Premise 1 of Vavova's formulation is subtly ambiguous. It can make a claim about what we *believe*, or what we *accept*.

(1-BELIEF) The correct response to peer disagreement is to stop *believing*.

(1-ACCEPTANCE) The correct response to peer disagreement is to stop *accepting*.

The first theory of acceptance denies (1-ACCEPTANCE) but is totally silent about (1-BELIEF). (1-BELIEF) is fully compatible with that theory. The conciliationists who argue for Vavova's Premise 1 intend to be arguing for (1-BELIEF). (The standard motivating cases – about check-splitting and so on – all involve belief, and not acceptance.) So the standard arguments for Premise 1 are not arguments against the first theory.

Now conciliationists also argue that peer disagreement should reduce our credence in the object of disagreement. And acceptance in moral contexts is plausibly tied to credences; it doesn't make sense to accept a proposition unless you have a higher credence in it than in the alternatives. But on the model introduced in §4, your credence in the metaphysically loaded part will already be quite low. Suppose that you're .9 in thirty particular moral beliefs. Then your credence that the ordering that captures those beliefs is the moral standard is strikingly low: .0424. But on the §4 model, you still accept it. You accept it because you're more confident in that proposition than in any alternative: your confidence in any alternative proposition about what's the moral standard is no higher than .00471. Given this conception of acceptance, reducing your credence in the metaphysically loaded part does not change what you accept, since acceptance only requires a *higher* credence in it than in the alternatives.

So the only way for a conciliationist to threaten the first theory would be to hold that you must distribute your credences *exactly equally* among all propositions that a peer disagrees with you about. Few conciliationists are that extreme – and for good reason! Consider two doctors reasoning independently: one ends up .97 that some particular dose will work, and the other ends up .95. The extreme conciliationist has to say that the first doctor should lower her credence that the dose will work. But that seems exactly backwards; it seems like further evidence that the first doctor is right!²⁵ The versions of conciliationism that are incompatible with the first theory seem too extreme to be plausible.

The second theory also predicts that peer disagreement is *not* a domain-specific reason against accepting the metaphysically loaded part. The second theory *broadens* the range of reasons that bear on accepting the metaphysically loaded part. It holds that our reactive attitudes can provide such domain-specific reasons – additional domain-specific reasons beyond those those given by coherence or incoherence with what we already accept. But it is *my* reactive attitudes that give me such reasons. So peer disagreement can only be relevant if it pushes me to abandon some of my particular reactive attitudes. And it's hard to see how that could happen. Neither a peer's lacking the reactive attitude nor a peer having an incompatible attitude seems to do.

Now a peer could give me *reasons* for abandoning my resentment; for example, she might point out that stepping on the toe was the only way to avoid much worse harm. But I'll abandon my resentment only if I already accept the normative judgments that my peer is appealing to: for example, the judgment that it's permissible to cause a small harm to avoid a larger

²⁵This example is from David Christensen (2009).

harm. In that case, I'd be abandoning my resentment because it was incoherent with judgments that I already accept, not because of peer disagreement. Bare peer disagreement doesn't itself ground reasons for abandoning resentment or other reactive attitudes. (At the very least, taking bare peer disagreement to provide such reasons is no part of the standard conciliationist package, so the NEW QUESTION still raises very different questions than the TRADITIONAL ONE.) So the second theory likely agrees with the first, that peer disagreement does not provide domain-specific reasons against accepting the metaphysically loaded part. The answer to the NEW QUESTION would then be 'no'.

Let's now suppose that the answer to the NEW QUESTION is 'no', and ask whether the argument from peer disagreement threatens moral knowledge. I say that it does not. Consider two disagreeing parties: A thinks that ϕ -ing is wrong, and B thinks that ϕ -ing is not wrong. Because they disagree, they accept different orderings to be the moral standard: A accepts one (\mathbf{o}_A) that forbids ϕ -ing, and B accepts one (\mathbf{o}_B) that doesn't. So if A knows that ϕ -ing is wrong, she knows that \mathbf{o}_A forbids ϕ -ing while accepting that \mathbf{o}_A is the moral standard. And if B knows that ϕ -ing isn't wrong, she knows that \mathbf{o}_B doesn't forbid ϕ -ing while accepting that \mathbf{o}_B is the moral standard. There isn't peer disagreement about what \mathbf{o}_A forbids, or peer disagreement about what \mathbf{o}_B forbids. So peer disagreement doesn't prevent the disagreeing parties from knowing the propositions they would need to know to have moral knowledge, because there *isn't* peer disagreement about those propositions. So peer disagreement can't undercut moral knowledge.

Let's unpack the crucial point more slowly. As noted earlier, Vavova's formulation of the argument is ambiguous.

(1-BELIEF) The correct response to peer disagreement *in belief* is to stop *believing*.

(1-ACCEPTANCE) The correct response to peer disagreement *in acceptance* is to stop *accepting*.

(2-BELIEF) There is a lot of peer disagreement *in beliefs* about morality.

(2-ACCEPTANCE) There is a lot of peer disagreement *in acceptance* about morality.

If the answer to the NEW QUESTION is 'no', then (2-BELIEF) is false.²⁶ There isn't peer disagreement about whether \mathbf{o}_A forbids ϕ -ing. Now (2-ACCEPTANCE) is still true; peers do disagree about whether \mathbf{o}_A is the moral standard. But given either of the two theories of acceptance, (1-ACCEPTANCE) is false. So given either of those theories, there is no sound

²⁶If the answer to the NEW QUESTION were 'yes', then A's belief would be the belief that \mathbf{o}_A is the moral standard and forbids ϕ -ing – and there would be peer disagreement about that belief, because there would be peer disagreement about what is the moral standard. But if the answer to the NEW QUESTION is 'no', the metaphysically loaded part, about \mathbf{o}_A being the moral standard, isn't locally accommodated, and the belief is just that \mathbf{o}_A forbids ϕ -ing.

inference to the conclusion that we should be agnostic about much of morality, because (1-ACCEPTANCE) and (2-BELIEF) are each false.

On the two theories of acceptance, then, the argument from peer disagreement is posing a pseudo-problem. It's forceful only if the norms on acceptance are the same as the norms on belief – and the norms on acceptance are different. At the same time, though, it's easy to see how we could have thought that this pseudo-problem is a genuine problem. Our ordinary beliefs (about check-splitting, or our physical environment) aren't hybrids of belief and acceptance; they're just non-hybrid beliefs. Though it's natural to generalize straightforwardly from our ordinary beliefs to our moral beliefs, it's a mistake to do so.

The overarching ambition of this paper has been to argue that moral epistemologists have been focused on too narrow a range of questions. They haven't focused on the NEW QUESTION, and the NEW QUESTION raises very different issues than those that moral epistemologists traditionally discuss. This paper has given decisive evidence that moral epistemologists need to address the NEW QUESTION before they can address their traditional questions; the NEW QUESTION is pressing when we realize that the metaphysically loaded part of moral discourse is a presupposition.

So my ambition has been to open up new vistas in moral epistemology. Moral epistemology must start with a theory of acceptance for the moral case: a theory of the domain-specific reasons for and against accepting the metaphysically loaded part of moral discourse. The Strawson-inspired approach strikes me as the most promising one, though there isn't space to argue that it is. This paper has worked out the conceptual foundations for this Strawsonian approach, explaining why conclusions about what we accept would have sweeping consequences for moral epistemology. In so doing, the paper has also worked out the conceptual foundations for a wide range of other possible approaches to moral epistemology; for example, it has also worked out the foundations for a very novel coherentist approach to moral epistemology. And other theories about acceptance in the moral case are possible too. But in each case, debates in moral epistemology should look very different once they incorporate the lessons from this paper.

References

- Abrusán, Márta. 2011. "Predicting the presuppositions of soft triggers." *Linguistics and Philosophy* 34:491–535.
- Amaral, P, Roberts, C, and Smith, E A. 2007. "Reivew of the Logic of Conventional Implicature by Chris Potts." *Linguistics & Philosophy* 30:707–749.
- Ayer, Alfred Jules. 1936. *Language, Truth and Logic*. London: Victor Gollancz Ltd.

- Brink, David. 1989. *Moral Realism and the Foundations of Ethics*. Cambridge: Cambridge University Press.
- Butler, Joseph. 1729. "Sermons." In L A Selby-Bigge (ed.), *British Moralists: Being Selections from Writers Principally of the Eighteenth Century*, 181–244. New York: The Bobbs-Merrill Company.
- Christensen, David. 2009. "Disagreement as Evidence: The Epistemology of Controversy." *Philosophy Compass* 4 (5):756–767.
- Crníc, Luka. 2011. *Getting \textit{Even}*. Ph.D. thesis, Linguistics and Philosophy.
- Daniels, Norman. 1979. "Wide Reflective Equilibrium and Theory Acceptance in Ethics." *The Journal of Philosophy* 76 (5):256–282.
- DeRose, Keith. 2005. "The Ordinary Language Basis for Contextualism, and the New Invariantism." *The Philosophical Quarterly* 55 (219):172–198.
- Elbourne, Paul. 2005. *Situations and Individuals*. The MIT Press: Cambridge, MA.
- Fine, Kit. 2009. *Semantic Relationalism*. Wiley-Blackwell.
- Finlay, Stephen. 2014. *Confusion of Tongues: A Theory of Normative Language*. Oxford University Press.
- Foot, Philippa. 2001. *Natural Goodness*. Oxford: Oxford University Press.
- Geurts, Bart. 1998. "Presuppositions and Anaphors in Attitude Contexts." *Linguistics and Philosophy* 21:545. ISSN 0165-0157.
- Harman, Gilbert. 1977. *The Nature of Morality*. Oxford University Press.
- Heim, Irene. 1992. "Presupposition Projection and the Semantics of Attitude Verbs." *Journal of Semantics* 9:183–221.
- Hume, David. 1739. *A Treatise of Human Nature*. Oxford: Clarendon, 2nd edition.
- Karttunen, Lauri. 1974. "Presupposition and Linguistic Context." *Theoretical Linguistics* 1:181 – 94.
- Kolodny, Niko and MacFarlane, John. 2010. "Ifs and Oughts." *Journal of Philosophy* 107 (3):115–143.
- Kornblith, Hilary. 2000. "The Contextualist Evasion of Epistemology." *Philosophical Issues* 10:24–32.
- 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.
- Lewis, David. 1970. “How to Define Theoretical Terms.” *Journal of Philosophy* 67:427–446.
- . 1975. “Languages and Language.” In Keith Gunderson (ed.), *Language, Mind and Knowledge*, volume 7 of *Minnesota Studies in the Philosophy of Science*, 3–35. Minneapolis: University of Minnesota Press.
- . 1983. “New Work for a Theory of Universals.” *Australasian Journal of Philosophy* 61:343–377.
- . 1984. “Putnam’s Paradox.” *Australasian Journal of Philosophy* 62:221–236.
- López de Sa, Dan. 2008. “Presuppositions of Commonality: An Indexical Relativist Account of Disagreement.” In Garca-Carpintero and Kolbel (eds.), *Relative Truth*, 297–310. Oxford University Press.
- . 2015. “Expressing Disagreement: A Presuppositional Indexical Contextualist Relativist Account.” *Erkenntnis* 80:153–165.
- Mandelkern, Matthew. 2017. *Coordination in conversation*. Ph.D. thesis, MIT.
- Olson, Jonas. 2014. *Moral error theory: History, critique, defence*. Oxford: Oxford University Press.
- Parfit, Derek. 2011. *On What Matters*, volume 2. Oxford: Oxford University Press.
- Plantinga, Alvin. 2000. *Warranted Christian Belief*. Oxford: Oxford University Press.
- Potts, Christopher. 2007. “The expressive dimension.” *Theoretical Linguistics* 33:165–198. ISSN 03014428. doi:10.1515/TL.2007.011.
- Rawls, John. 1971. *A Theory of Justice*. Cambridge, MA: Harvard University Press.
- . 1975. “The Independence of Moral Theory.” *Proceedings and Addresses of the American Philosophical Association* 48:5–22.
- Roberts, Craige. ????. “Only, Presupposition and Implicature.” .
- . 2012. “Information Structure: Towards an integrated formal theory of pragmatics.” *Semantics and Pragmatics* 6:1–69.

- Rubinstein, Aynat. 2014. "On necessity and comparison." *Pacific Philosophical Quarterly* 95:512–554.
- . 2017. "Straddling the line between attitude verbs and necessity modals." In Ana Arregui, María Luisa Rivero, and Andrés Salanova (eds.), *Modality across syntactic categories*, 109–131. Oxford University Press.
- Scanlon, T. M. 1998. *What We Owe to Each Other*. Cambridge, MA: Harvard University Press.
- . 2002. "Rawls on Justification." In Samuel Freeman (ed.), *The Cambridge Companion to Rawls*, 139–167. Cambridge University Press.
- . 2014. *Being Realistic about Reasons*. New York: Oxford University Press.
- Schlenker, Philippe. 2003. "A Plea for Monsters." *Linguistics and Philosophy* 26:29–120. ISSN 0165-0157. doi:10.1023/A:1022225203544.
- . 2009. "Local contexts." *Semantics and Pragmatics* 2:1–78.
- Schroeder, Mark. 2007. *Slaves of the Passions*. Oxford: Oxford University Press.
- Setiya, Kieran. 2012. *Knowing Right From Wrong*. Oxford University Press.
- Sextus Empiricus. 2000. *Outlines of Scepticism*. Cambridge: Cambridge University Press.
- Silk, Alex. 2016. *Discourse Contextualism: A Framework for Contextualist Semantics and Pragmatics*. New York: Oxford University Press.
- . 2017a. "How to Embed an Epistemic Modal: Attitude Problems and Other Defects of Character." *Philosophical Studies* 174 (7):1773–1799.
- . 2017b. "Normative Language in Context." In Russ Shafer-Landau (ed.), *Oxford Studies in Metaethics, Volume 12*. Oxford University Press.
- Simons, Mandy, Beaver, David, Roberts, Craige, and Tonhauser, Judith. 2017. "The Best Question: Explaining the Projection Behavior of Factives." In Anton Benz and Katja Jasinskaja (eds.), *Discourse Processes*.
- Sinnott-Armstrong, Walter. 2006. *Moral skepticisms*. Oxford: Oxford University Press.
- Slooman, Aaron. 1970. "'Ought' and 'Better'." *Mind* 79 (315):385–394.
- Soames, Scott. 2009. "Presupposition." In *Philosophical Essays, Volume 1: Natural Language: What It Means and How We Use It*, 73–130. Princeton, NJ: Princeton University Press.
- Stalnaker, Robert. 1973. "Presuppositions." *Journal of Philosophical Logic* 2:447–457.

- . 2002. “Common Ground.” *Linguistics and Philosophy* 25:701–721.
- Strawson, P F. 1974. “Freedom and Resentment.” In *Freedom and Resentment and other essays*. Methuen and Co, Ltd.
- Street, Sharon. 2006. “A Darwinian Dilemma for Realist Theories of Value.” *Philosophical Studies* 127 (1):109–166.
- van Inwagen, Peter. 1996. “It is Wrong, Always, Everywhere, and for Anyone, to Believe Anything, Upon Insufficient Evidence.” In J. Jordan and D. Howard-Snyder (ed.), *Faith, Freedom, and Rationality*, 137–154. Hanham, MD: Rowman and Littlefield.
- Vavova, Katia. 2014. “Moral Disagreement and Moral Skepticism.” *Philosophical Perspectives* 28(1):302–333.
- 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.
- Villalta, Elisabeth. 2009. “Mood and gradability: An investigation of the subjunctive mood in Spanish.” *Linguistics and Philosophy* 31 (4):467–522.
- von Fintel, Kai. 1999. “NPI licensing, Strawson entailment, and context dependency.” *Journal of Semantics* 16 (2):97–148.
- . 2008. “What is presupposition accommodation again?” *Philosophical Perspectives* 22 (1):137–170.
- Wallace, R. Jay. 1998. *Responsibility and the Moral Sentiments*. Cambridge, MA: Harvard University Press.
- Willer, Malte. 2014. “Dynamic Thoughts on Ifs and Oughts.” *Philosophers’ Imprint* 14 (28):1–30.