Moral Vagueness as Semantic Vagueness*

Rohan Sud

Does moral vagueness require ontic vagueness? A central challenge for nonontic treatments of moral vagueness arises from the referential stability of moral terms across small changes in how they are applied: if moral vagueness is not ontic vagueness, it's hard to explain this referential stability. Pointing to this challenge, Miriam Schoenfield has argued that moral vagueness is ontic vagueness, at least for a moral realist. I disagree. I argue that a moral realist can use a conceptual role semantics for moral terms to give a purely semantic treatment of moral vagueness.

Consider the following case of moral vagueness, from Miriam Schoenfield:[1]

> **Amputations**: It is impermissible to amputate a person's arm to save another's life. It is permissible to amputate a person's arm to save a billion lives. How many lives must be at stake for it to be permissible to amputate someone's arm? Plausibly, we can create a Sorites series, admitting of borderline cases of permissibility, out of a series of amputations, each of which is performed to save an increasing number of lives.[2]

A sorites series of hair loss makes it plausible that some people are borderline bald. Similarly, sorites series like those described in Amputations make it plausible that some actions are borderline wrong. In gen-

* Special thanks to David Manley for helpful discussion. Thanks also to two anonymous referees, whose comments significantly improved the article.

1. Miriam Schoenfield, "Moral Vagueness Is Ontic Vagueness," *Ethics* 126 (January 2016): 257–82.

2. Ibid., 263.

Ethics 129 (July 2019): 684–705

eral, there are two theories of borderline or indeterminate cases.[3] Very roughly speaking, *semantic* theories of vagueness locate the indeterminacy in our language, rather than in the world being represented by the language, whereas *ontic* theories of vagueness locate the indeterminacy in the world being represented.

Most philosophers accept a semantic theory of indeterminate cases of baldness. A natural thought, then, is that this theory can be seamlessly extended from indeterminate cases of baldness to indeterminate cases of wrongness. After all, initial reflection on the sorites series of hair loss and the sorites series of amputations makes baldness and wrongness appear analogous in terms of their vagueness.

As Schoenfield and others have pointed out, however, the cases are importantly different.[4] Moral terms like 'wrong' are *stable*, in the sense that they would have the same application conditions even if our dispositions to apply the predicates were slightly different.[5] Consider moral twin earth cases.[6] Imagine a "twin" linguistic community that spoke an analogue of English, in which the term 'wrong' played the same practical role (e.g., members of this twin community blamed those they said, in their words, "did something that was wrong") but they applied the term to a slightly different class of actions. (Say, e.g., our twins were disposed toward a more lax usage of their term 'wrong' when confronted with Amputations than actual English speakers in fact are.) Intuitively, this twin community's term 'wrong' has the same application conditions as our homophonous term—despite their more lax usage. This stability, however, is not characteristic of typical nonmoral vague terms like 'bald'. If we imagine a linguistic community that is disposed toward a more lax usage of 'bald', intuitively their term 'bald' has more generous application conditions (even if there also remain cases of indeterminacy).

However, it is far from obvious whether accounts of semantic vagueness can be extended to stable terms. That's because it is far from obvi-

3. I am setting aside epistemic theories of indeterminacy such as that in Timothy Williamson, *Vagueness* (London: Routledge, 1994). (Schoenfield gives such views extensive consideration.) Strictly speaking, there may be a difference between cases of indeterminacy (or borderline cases) and cases of vagueness. For instance, vagueness might require indeterminacy and sorites susceptibility. In the first instance, then, we're investigating cases of moral indeterminacy, but I will sometimes gloss over this difference and use the terms 'indeterminate', 'borderline', and 'vague' interchangeably.

4. A similar puzzle is raised in David Manley, "Moral Realism and Semantic Plasticity" (unpublished manuscript).

5. Schoenfield uses the term 'rigid', but I'll use the term 'stable' in order to prevent confusion with Kripke's distinct notion of rigidity.

6. Terry Horgan and Mark Timmons, "Troubles for New Wave Moral Semantics: The Open Question Argument Revived," *Philosophical Papers* 21 (1992): 153–75. Schoenfield, "Moral Vagueness," 265–66, contains a rich and novel argument for stability which I cannot discuss here.

ous that a metasemantic theory can simultaneously predict both that moral terms are stable and that moral vagueness is semantic. The curious combination of vagueness with stability exhibited by wrongness thus raises a challenge:

> **Challenge**: Find a metasemantic theory of reference that predicts that moral terms are stable and that predicts that moral vagueness is purely semantic rather than ontic.

Reflecting on the vague yet stable feature of wrongness, Schoenfield concludes that insofar as we are moral realists—that is, insofar as we accept that moral properties are fundamental—the Challenge cannot be met.[7] Because she accepts that moral terms are stable, she concludes that moral vagueness is ontic vagueness:[8]

> **MVIOV**: If moral properties are fundamental, then there is ontic vagueness.[9]

In this article, I will attempt to meet the Challenge and rebut Schoenfield's argument for MVIOV. Borrowing from work by Ralph Wedgwood and J. R. G. Williams,[10] I will show that conceptual role metasemantic theories (CRS) are up to the Challenge: they predict that moral terms are stable and semantically, rather than ontically, vague. And because my explanation is open to those who think that moral properties are fundamental, I'll have rebutted Schoenfield's argument for MVIOV.

It's worth noting that, because the structure of the present argument is so general, our discussion has a significance far beyond the moral domain. Expressions like 'being conscious', 'being a person', and 'composing a new object' are plausibly stable and susceptible to sorites sequences. If the combination of stability and vagueness supports ontic vagueness, then we would have a new path from vagueness in personal identity, consciousness, and mereological composition to ontic vagueness. Those of us who are opposed to ontic vagueness, then, have strong reason to pay attention to the present Challenge.

7. Manley ("Moral Realism") also takes semantic theories of vagueness to be in tension with the stability of moral terms.

8. Schoenfield also argues that explanations of moral vagueness cannot be epistemic. See n. 3.

9. Officially, Schoenfield remains neutral as to whether there is moral vagueness and includes the claim that there is moral vagueness in the antecedent of the conditional. However, she suggests that there is strong reason to think that there is moral vagueness. I agree and will assume so throughout.

10. Ralph Wedgwood, "Conceptual Role Semantics for Moral Terms," *Philosophical Review* 110 (2001): 1–30; J. R. G. Williams, "Normative Reference Magnets," *Philosophical Review* 127 (2018): 41–71.

I. BASIC COMMITMENTS

I will be arguing that the stability of moral terms does not support MVIOV. That is, I will show that the stability of moral terms creates no special problem for a realist who thinks that moral vagueness is semantic. But I want to first note a basic commitment of any moral realist who thinks that moral vagueness is semantic rather than ontic—a commitment that isn't related to stability and one we should be willing to take on board for the purpose of evaluating the Challenge-based argument for MVIOV.

The basic commitment I have in mind is to a plenitude of precise moral properties. To see why this is a commitment of any moral realist who thinks that moral vagueness is semantic, note that semantic vagueness requires that it is indeterminate which of a plurality of referents our terms refer to.[11] Consider a mundane case of vagueness:

(1) Harry is such that it's indeterminate whether he is bald.

If the vagueness involved in (1) is purely semantic rather than metaphysical, then there are a plenitude of "precise" properties (properties like *having fewer than 49,999 hairs*, *having fewer than 50,000 hairs*, etc.) such that it's indeterminate which one of these precise properties 'bald' refers to.[12]

Applied to a claim of moral indeterminacy, such as

(2) Act *a* is such that it's indeterminate whether it is wrong,

semantic moral vagueness requires that there are a plenitude of precise moral properties (properties like wrong_1, wrong_2, etc.) such that it's indeterminate which of these precise properties moral terms in English (like 'wrong') refer to.

Given these plenitudinous precise moral properties corresponding to vague English moral predicates, there are two options for the moral realist:

11. Strictly speaking, a proponent of a nonstandard treatment of semantic vagueness might deny this. In other work, for instance, I have claimed that vagueness results from our tokening a plenitude of (determinately referring) terms rather than a vaguely referring term. Although I will assume the more standard treatments here, I don't see any obstacle for reformulating the present discussion in a way that conforms with this nonstandard treatment. See Rohan Sud, "Plurivaluationism, Supersententialism, and the Problem of the Many Languages," *Synthese*, forthcoming.

12. A property p is precise (in this sense) if and only if nothing is such that it indeterminately instantiates p. The locution "it's indeterminate which of some *X*s is $\phi(x)$" is short for the claim "for each of the *X*s, it's indeterminate whether it is $\phi(x)$, but exactly one of the *X*s is $\phi(x)$."

Plenitudinous Fundamentality: The plenitude of precise moral properties are fundamental.

Sparse Fundamentality: The plenitude of precise moral properties are nonfundamental but hold in virtue of some yet-to-be-discovered fundamental moral properties described in some yet-to-be-discovered completed moral theory.

We might not find either option particularly attractive. But the fine-grainedness of moral reality is a basic commitment for any moral realist who thinks that moral vagueness is semantic rather that ontic. For the purposes of assessing the Challenge-based argument, then, we must bracket our concerns with this commitment and provisionally accept that there are a plenitude of precise referents for 'wrong' that either are fundamental or hold in virtue of a yet-to-be-discovered completed moral theory. It's up to a fan of the Challenge-based argument to show why the stability of moral terms should lead us to reduce our initial credence in this claim or else to show that stability leads to some (other?) unsavory consequence for the realist who rejects ontic vagueness.

II. CLARIFYING THE CHALLENGE

Our Challenge is to give a metasemantic theory that predicts that moral vagueness is semantic vagueness and that moral terms are stable. In this section, we will clarify that Challenge.

A. Vagueness

In order to describe how a metasemantic theory might predict that moral vagueness is semantic rather than ontic, we first need to regiment our rough characterizations of semantic and ontic vagueness. Say that a 'fundamental language' is a language with predicates that refer to all and only fundamental properties.[13] When Schoenfield lays down the Challenge, she characterizes ontic vagueness as vagueness "that would remain even if we spoke a [fundamental] language."[14] As I understand it, vague-

13. I use the expression 'fundamental language' where Schoenfield uses the expression 'perfect language', which she defines as "a language that contains all and only predicates that are necessary to provide a complete and accurate description of how things are fundamentally" ("Moral Vagueness," 260). I prefer my own terminology because the term 'perfect' suggests that such languages are referentially determinate, but, as I'll argue, such languages can be referentially indeterminate and so in that sense less than perfect.

14. Schoenfield is following Elizabeth Barnes, "Fundamental Indeterminacy," *Analytic Philosophy* 55 (2014): 339–62. (However, I don't think she follows her closely enough; see nn. 17 and 18.) She also adds the requirement that the vagueness remains for omniscient speakers, which is meant to rule out epistemic treatments of vagueness. Because we're ignoring such treatments (see n. 3), I've elided this requirement.

ness "remains in a language" when some claim in that language is indeterminate.[15] So, according to Schoenfield, we have the following:

> **Ontic Vagueness***: There is ontic vagueness just in case, were we to speak a language that referred to all and only fundamental properties, a claim in that language would be indeterminate.

Schoenfield claims that her characterization "gets at the phenomenon" we're supposed to be interested in.[16] I think that's roughly right, but we need to make a couple of minor clarifications to Schoenfield's understanding of ontic vagueness that help us get at the interesting notion. I think that these clarifications should be relatively uncontroversial (but are relevant for what will follow).

First, we shouldn't take too seriously the imagery of us speaking a fundamental language.[17] If a fundamental language cannot be spoken by creatures like us (because, e.g., such a language contains more predicates than we are able to learn), that does not trivially entail that there is ontic vagueness—at least on any interesting sense of 'ontic vagueness'. More generally, a fundamental language need not be a natural language at all. Instead, it may be an artificial language. As such, we should think of such a fundamental language in the abstract, as simply a pairing of a particular syntax and a particular semantic model that maps terms with referents, even if such a language is not possibly spoken. As such, the way the terms in this fundamental language get their meaning need not be the way the terms in a natural language like English get their meaning.

Second, we should distinguish between *referentially determinate languages*, in which all of the terms have their referents determinately (i.e., for all x, each term in the language determinately refers to x or determinately fails to refer to x), and *referentially indeterminate languages*, in which it's indeterminate what some terms refer to.[18] We should restrict our attention to referentially determinate languages. Otherwise, we risk trivializing the the-

15. More specifically, vagueness remains when a claim of the form "it's indeterminate that ϕ" is true in that language (when that language is supplemented with an indeterminacy operator).

16. Schoenfield, "Moral Vagueness," 260.

17. Although Barnes ("Fundamental Indeterminacy") initially introduces the notion of metaphysical indeterminacy using the imagery of a spoken fundamental language, she quickly drops it. When she's being more careful, she writes, "A theory counts as committed to fundamental indeterminacy just in case the basic/fundamental/most natural/etc. description of that theory includes sentences which are indeterminate" (ibid., 347).

18. In her earlier work, Barnes makes clear that ontic vagueness is supposed to be vagueness that remains in a language whose representational content is precisified. Elizabeth Barnes, "Ontic Vagueness: A Guide for the Perplexed," *Noûs* 44 (2010): 601–27. And in a later article, she makes clear that ontic vagueness must not "arise from imprecision in language" ("Fundamental Indeterminacy," 358; see also her n. 38).

sis of ontic vagueness. Here's why. Imagine a language with terms 'elecuark' and 'quarton' where we make the following stipulation: the terms are to refer to one of the fundamental properties *being an electron* or *being a quark*, but they are to refer to different properties. We don't say anything about which property they refer to. This language is a fundamental language, but it's indeterminate which fundamental property the predicates 'elecuark' and 'quarton' refer to. So, a claim in this referentially indeterminate language will be indeterminate. Surely this cannot be enough for ontic vagueness—at least on any interesting sense of 'ontic vagueness'! After all, this vagueness would be due to semantic vagueness (vagueness in which referents our words refer to) rather than vagueness in the world itself (vagueness in the referents themselves).

Dropping the presupposition that we are able to speak fundamental languages and restricting attention to referentially determinate languages, we can restate our characterization of ontic vagueness as follows:

> **Ontic Vagueness**: There is ontic vagueness just in case there is a referentially determinate language that refers to all and only fundamental properties and a claim in that language is indeterminate.

Let's take Ontic Vagueness as our official characterization of ontic vagueness.[19]

Recall that if vagueness is semantic rather than ontic, we must posit vagueness as to what our terms refer to.[20] In particular, if there is no ontic vagueness, then the prejacents in (1) and (2) must be stated in referentially indeterminate language. The obvious candidates are 'bald' and 'wrong'. So it must be indeterminate which of a bunch of precise properties 'bald' and 'wrong' refer to.

We can restate this claim using the notion of a *candidate reference relation*. If it is vague what our terms refer to, then there are several different "precise" relations between the terms in our language and potential referents, such that it's indeterminate which of those relations coin-

19. Someone might worry that Ontic Vagueness is not broad enough to capture the intuitive notion of ontic vagueness owing to its focus on the fundamental. Instead, he might prefer that we characterize ontic vagueness more broadly as vagueness that remains in any referentially determinate language—regardless of whether its subject matter is the fundamental or the nonfundamental. (In this way, it's more similar to the interpretation of metaphysical indeterminacy in Barnes, "Ontic Vagueness.") After all, if there is vagueness in the nonfundamental world, there is still vagueness in the world. Although we'll stick with Ontic Vagueness as our official regimentation, I am sympathetic to this concern: our solution for the realist who rejects ontic vagueness applies equally well when the notion is given this broader characterization.

20. This follows immediately from the broader characterization of ontic vagueness suggested in n. 19.

cides with the relation of reference. Let's call these relations "candidate reference relations." So, if there is no ontic vagueness and it is vague what 'bald' refers to, one candidate reference relation relates 'bald' with the property of *having fewer than 50,000 hairs*. A different one relates 'bald' with the property of *having fewer than 49,999 hairs*. And so forth. Similarly with 'wrong'. So, we can restate the claim that it is vague what our terms refer to as the claim that there are multiple candidate reference relations.

If there is no ontic vagueness, it is vague what our terms refer to (i.e., there are multiple candidate reference relations). But our metasemantic theories make predictions about the reference of our terms—including whether it is vague what the terms refer to. This establishes a link between our metasemantic commitments and whether there is ontic vagueness: if there is no ontic vagueness, our metasemantic commitments must make room for vague reference. How do they make this room?

Metasemantic theories tell us what criteria (dispositions to use, naturalness, rationality, causation, etc.) fix reference. As such, these theories contain claims of the form

> **Schematic Metasemantic Theory**: The reference relation is the relation *r* between words and referents that meets such and such criteria,

where different metasemantic theories of reference differ with respect to how they fill in the reference-fixing criteria. For example, one way of filling in the above schema gives us the widely discussed metasemantic theory "naive reference magnetism" (NRM):[21]

> **NRM**: The reference relation is the relation *r* between words and referents that maximizes the combined degree of fit and eligibility, where a relation "fits with usage" to the degree that it makes true the sentences the linguistic community is strongly disposed to accept and "is eligible" to the degree that it assigns more natural (in the Lewisian sense) referents to words.

Our metasemantic theory gives us the reference-fixing criteria. But if the criteria fail to privilege a unique relation between our words and potential referents, reference will be vague. One way the criteria might fail to privilege a unique relation is if there are multiple relations that meet the criteria—that is, multiple relations are tied with respect to

21. See, among others, David Lewis, "Putnam's Paradox," *Australasian Journal of Philosophy* 62 (1984): 221–36; David Lewis, "New Work for a Theory of Universals," *Australasian Journal of Philosophy* 61 (1983): 343–77; Theodore Sider, "Criteria of Personal Identity and the Limits of Conceptual Analysis," *Philosophical Perspectives* 15 (2001): 189–209; Brian Weatherson, "What Good Are Counterexamples?," *Philosophical Studies* 115 (2003): 1–31.

the reference-fixing criteria. In this case, each relation is a candidate reference relation and we have vague reference. For instance, suppose that we accept NRM and there are multiple relations that maximize the combined degree of fit and eligibility.[22] Then, we'd expect it to be vague what our terms refer to. When this is the case, let's say that a metasemantic theory predicts vague reference *by ties*.

B. Stability

As we've seen, if a metasemantic theory is to predict that moral vagueness is semantic, it must predict that English moral terms vaguely refer. Now, let's consider the other horn of our Challenge: how does a metasemantic theory predict stability?

To keep things simple, first consider terms that determinately, rather than vaguely, refer. Such a term is stable when its "twin" term (the term used by our counterfactual twins who apply the term to a slightly different class of objects) has the same referent. But the metasemantic theory tells us how terms refer across such communities. As such, metasemantic theories make certain predictions of stability.

One way a metasemantic theory can predict stability is via what I'll call the *trumping* strategy. According to the trumping strategy, our disposition to apply a term to a particular class of objects is one factor that fixes reference, but that factor can be trumped by other competing factors. Stability in a term results when these other factors trump shifts in our dispositions to apply the term.

Continuing with our example of NRM, a proponent of this metasemantic theory can deploy the trumping strategy in order to predict the stability of terms that refer to highly eligible properties such as *being an electron*. In a counterfactual scenario in which we are disposed to apply the term 'electron' to a slightly different class of objects, a slightly different assignment of reference would have a higher degree of fit than the actual assignment. So, this shift in dispositions would exert some metasemantic pressure to shift reference. However, NRM can predict the stability of 'electron' in the face of this metasemantic pressure if *being an electron* is a particularly natural property. The actual assignment of reference would have a higher degree of eligibility than the shifted assignment. Because reference is determined by the combined degree of fit, the particular naturalness of our actual referents exerts a countervailing metasemantic pressure to hold reference fixed. If the actual referents are much more natural than those assigned by the rival shifted assignment and degree of eligibility is weighted heavily when degree of fit and eligibility are combined, then the pressure to shift reference exerted by our shift

22. See Schoenfield, "Moral Vagueness," 270, who takes the suggestion from Tom Dougherty, "Vague Value," *Philosophy and Phenomenological Research* 89 (2014): 352–72.

in our dispositions to apply the terms will be trumped by the countervailing pressure exerted by the eligibility of the actual assignment of reference.[23]

Now let's consider terms that vaguely refer, as a proponent of semantic vagueness believes is the case with terms like 'bald' and 'wrong'. For terms that vaguely refer, the relevant notion of stability is a bit stronger than for terms that determinately refer. Such a term is stable when it has all of the same candidate reference relations across small changes in the community's use of the term. I take it that moral terms have this stronger notion of stability. For in moral twin earth, our twins' moral terms determinately have the same referent as ours—our moral talk determinately has the same subject matter.

C. Semantic Vagueness with Stability: The Problem

We are now in a position to see why the Challenge is so challenging. If there is no ontic vagueness, it's indeterminate which of several properties 'wrong' refers to. That is, there are many candidate reference relations, none of which are uniquely privileged by the reference-fixing criteria given by our metasemantic theory. And if 'wrong' is stable, the reference-fixing criteria must delimit these same candidate reference relations across shifts in usage. But it's not obvious how this can work. For our change in usage may change the class of candidate reference relations.

Consider again NRM. Assuming that NRM predicts vague reference of moral terms by ties, there are multiple precise relations between our words and referents that each maximize combined degree of fit and eligibility. But if our usage shifted a bit, a different range of candidate reference relations would be delimited. This is most obvious for a twin community with application dispositions that privilege a proper subset of the actual candidate reference relations. Given that the actual candidate reference relations each maximized combined degree of fit and eligibility for the actual linguistic community and assuming that eligibility doesn't vary across communities, if the counterfactual community's usage bumps up the degree of fit for some proper subset of these relations—even ever so slightly—these bumped-up relations will now have a higher combined degree of fit and eligibility than the other actual candidate reference relations. So, the term would not be stable. Even if use is trumped by eligibility, deviations in use will still break the ties that lead to vagueness.[24]

23. See Schoenfield, "Moral Vagueness," 267. For a detailed discussion, not in the context of vagueness, of reference magnetism and stability for moral terms, see Billy Dunaway and Tristram McPherson, "Reference Magnetism as a Solution to the Moral Twin Earth Problem," *Ergo* 3 (2016): 639–79.

24. Schoenfield, "Moral Vagueness," 270–71.

We can sum up the problem as follows. Moral terms are stable; our strategy for explaining stability was the trumping strategy. Rejecting ontic vagueness requires moral terms to vaguely refer; our strategy for explaining vague reference was via ties in the reference-fixing criteria. But the trumping strategy is in tension with vagueness predicted by such ties.

III. THE SOLUTION: CRS FOR MORAL TERMS

The way to meet the Challenge is to adopt a conceptual role metasemantic (CRS) theory.

A. CRS for Moral Terms

According to (one version of[25]) CRS, reference of our terms is fixed by the theoretical and practical inferences we make with sentences involving that term.[26] More specifically, our sentences are associated with various "core" theoretical and practical inferences, from which we infer sentences from sentences or infer actions from sentences. And, according to CRS:

> **CRS**: The reference relation is the relation *r* between words and referents that assigns to words the referents that determine contents of the sentences of the inference which maximizes the rationality of the core inferences involving those sentences.

Following Wedgwood, Williams, and others, we can apply CRS to moral terms.[27] Our word 'wrong' is associated with certain practical—rather than theoretical—inferences from sentences to actions, and the term refers to whatever most rationalizes these practical inferences. For example, from sentences like "*a* is wrong," I infer that certain reactive attitudes (e.g., blame) are to be taken toward those that perform *a*. Ref-

25. Of course, there are a number of versions of CRS, even as applied to moral terms. Here I'm focused on just one version of those theories. However, for presentational purposes, I will continue to use the more general term 'CRS' for this particular version.

26. Officially, the sentences that figure in the practical and theoretical inferences are, in the first instance, sentences tokened "in our belief box" in a language of thought. It's these sentences of "mentalese" that are assigned reference in the way described in the main text—sentences of spoken and written language are assigned reference derivatively, based on the reference of the mentalese sentences that correspond with the spoken and written sentences of the communal language. For ease of presentation, I simply talk as if the sentences that figure in the inferences are sentences in the spoken and written language.

27. Wedgwood, "Conceptual Role Semantics"; Williams, "Normative Reference Magnets." My presentation here and in the next subsection follows Williams most closely (e.g., I take the conceptual role of normative terms to be primarily tied to our reactive attitudes), although I do not claim to follow him exactly. Part of why I follow Williams is because he explains very clearly how CRS can predict stability. (Schoenfield follows Wedgwood.)

erence is fixed in a way that makes these practical inferences most rational.[28] This theory makes use of a claim about our reasons: it is most rational to make these practical inferences from a sentence that means that *a* is wrong. That is, it is less rational to move from a sentence that means that Donald Trump disapproves of *a* to the action of blaming those that perform *a*. With this normative assumption in hand, CRS predicts that 'wrong' refers to *wrongness* rather than *disapproved of by Donald Trump*.

B. CRS Explains Stability

Above, we saw that NRM deploys the trumping strategy to explain stability. There is a second strategy, in addition to the trumping strategy for predicting stability—what I'll call the *impotence* strategy.[29] On this strategy, our disposition to apply a term to a particular class of objects is not a factor that fixes reference, even if other, more nuanced features of our use are. Because our application dispositions are referentially impotent, stability is a direct result.

CRS (at least on one way of developing the view) deploys the impotence strategy: our disposition to apply the term 'wrong' to this or that class of actions is referentially impotent. Rather, the relevant reference-fixing features of our use are exhausted by the core practical inferences from sentences that include 'wrong'. These practical inferences—and the facts about which practical inferences are most rational—do not vary across the counterfactual scenarios in which we most clearly intuit stability. That is, the twin moral term plays the same practical role in the twin community. As such, stability is a direct prediction: even if our twins are disposed to apply 'wrong' to a slightly different class of actions, assigning the twin term 'wrong' to wrongness still maximizes the rationality of the shared practical inferences.[30]

C. CRS Explains Vagueness

According to CRS, the reference-fixing criteria for a term *t* are given by the rationality of the core inferences involving that term. If those criteria don't privilege a unique referent, there will be semantic vagueness. In the case of our term 'wrong', the reference-fixing criterion is given by the rationality of the practical inferences from the sentences "*a* is wrong" to (inter alia) the action of blaming those that do *a*—'wrong' refers to

28. The relevant notion of rationality here is supposed to be substantive rationality rather than structural rationality—see secs. 2.1 and 2.2 of Williams, "Normative Reference Magnets."

29. Although Schoenfield doesn't explicitly distinguish between the two strategies, she appears to suggest the impotence strategy to proponents of CRS; see "Moral Vagueness," 267.

30. For more nuanced discussion on the relationship between CRS and stability, see Williams, "Normative Reference Magnets."

whatever referent makes this inference most rational. In this case, how might this reference-fixing criterion fail to privilege a unique referent?

We've already discussed, in the context of NRM, one way the reference-fixing criteria might fail to privilege a unique referent: multiple relations can be tied with respect to the reference-fixing criteria. In such a case, each of the relations is a candidate reference relation and there is vague reference. Applied to the present context, CRS predicts that 'wrong' is vague if there are multiple properties that make our practical inference most rational.[31] In this case, there are multiple candidate reference relations relating our word 'wrong' to these multiple most rationalizing properties, and there is vague reference.

As Schoenfield recognizes, in addition to predicting semantic vagueness via ties, there is a second, underappreciated way in which the reference-fixing criteria can fail to privilege a unique referent: it can be indeterminate which of several relations the reference-fixing criteria apply to. In this case, each such relation is a candidate reference relation. When this happens, let's say that the metasemantic theory predicts semantic vagueness via *indeterminate satisfaction*.

Consider the case of NRM. As David Lewis notes in passing,[32] NRM predicts semantic vagueness via indeterminate satisfaction if it's indeterminate which particular relation maximizes combined degree of fit and eligibility—perhaps owing to indeterminacy in how to weigh fit and eligibility. And a simple Kripkean theory of reference predicts that the term 'Kilimanjaro' will be vague if it's indeterminate which precise hunk of rock was demonstrated in the initial dubbing ceremony. Predicting vague reference via indeterminate satisfaction does not require ontic vagueness (or at least not obviously so) if the metasemantic theory deploys vague language in stating the reference-fixing criteria. For then it can be semantically indeterminate whether the criteria apply to this or that relation.[33] And when we consider the terms used to state the criteria—expressions

31. Interestingly, Schoenfield doesn't consider the possibility of CRS predicting vague reference via ties (although she does in the case of NRM). She writes, "If it can be semantically indeterminate which of a pair of actions is better, this must be because, either it's indeterminate which inference rules are the ones that constitute the conceptual role for 'better than,' or it's indeterminate which relation makes those rules valid" ("Moral Vagueness," 272). Neither option corresponds to multiple relations (determinately) making the rules (that determinately constitute the conceptual role) valid.

32. After sketching his theory of NRM, David Lewis writes, "The terms of trade [between truth of theory and eligibility] are vague; that will make for moderate indeterminacy of reference" ("Putnam's Paradox," 67).

33. For more on indeterminate satisfaction in the context of NRM, see Rohan Sud, "Vague Naturalness as Ersatz Metaphysical Vagueness," in *Oxford Studies in Metaphysics*, ed. Karen Bennett and Dean W. Zimmerman (Oxford: Oxford University Press, 2018), 11:243–77. In that article, I suggest that the term 'perfectly natural' may be semantically vague, in which case NRM predicts various sorts of semantic vagueness via indeterminate satisfaction.

like 'dispositions to use', 'combined degree', 'causation', etc.—it's plausible that this language is vague. So, it's plausibly semantically indeterminate which relation meets the criteria stated in these terms.

Applied to the case of CRS for moral terms, if it's indeterminate which of some properties most rationalizes the practical inference from the sentences "*a* is wrong" to the action of, say, blaming those that do *a*, it will be indeterminate which of these properties the term 'wrong' refers to. And CRS can predict vague reference of moral terms via indeterminate satisfaction without obviously involving ontic vagueness if the reference-fixing criterion is stated in vague language, for then it can be semantically indeterminate which referent most rationalizes the inferences. The reference-fixing criterion may be indeterminate, for instance, because the notion of blame is semantically indeterminate, ripe for precisification. Or the reference-fixing criterion may be indeterminate, for instance, because the notion of rationality is semantically indeterminate, also ripe for precisification.[34]

Summing up, there are two ways CRS can predict that moral terms are vague (i.e., that there are multiple candidate reference relations). CRS can predict that moral terms are vague via ties if there are multiple properties that most rationalize the relevant practical inferences. And CRS can predict that moral terms are vague via indeterminate satisfaction if it's indeterminate which of multiple properties most rationalizes the relevant practical inferences. The latter route doesn't obviously involve ontic vagueness if the reference-fixing criterion is stated in vague language that can be made precise—for instance, if the notion of blame is semantically vague.

We've seen two ways that CRS can predict that moral terms like 'wrong' are referentially vague. And, upon reflection, it's quite plausible that CRS does predict that moral terms like 'wrong' are referentially vague. In order to illustrate the plausibility of this claim, we'll freely assume that (instead of some "vague property" that is vaguely instantiated by various actions) there is a range of "precise" moral properties—$\text{wrong}_1, \ldots, \text{wrong}_n$—such that in borderline cases an action has some but not all of these properties. You can think of each of these precise normative properties as corresponding to sharp divisions we might draw to resolve borderline cases in a moral sorites sequence like Amputations. (If you feel uncomfortable with us assuming this plenitude of fine-grained moral properties, recall the discussion in Sec. I.) To simplify

34. Assuming that indeterminacy about rationality is in fact semantic doesn't beg the question against Schoenfield. It would beg the question for our opponent to assume that normative indeterminacy (including rational indeterminacy) must be ontic when making the case that moral indeterminacy must be ontic. But we are merely showing that there is a consistent package of view on which moral indeterminacy is semantic, so we can include in that package the claim that rational indeterminacy is semantic.

matters, we can limit our consideration to two such properties: wrong_1 and wrong_2. Suppose, for example, that it's indeterminate whether it's wrong to amputate an innocent bystander's arm to save five people's lives. Then, that action has one precise normative property (wrong_1) but lacks another (wrong_2).

With these properties in hand, consider again the quintessential practical inferences that fix the reference of 'wrong' according to CRS. From sentences like '*a* is wrong', I infer (inter alia) that certain reactive attitudes (e.g., blame) are to be taken toward those that perform *a*. Now ask yourself, which property, when assigned as the referent of 'wrong', makes these practical inferences most rational? Is it most rational to draw these practical inferences from the sentence with the content that *a* is wrong_1 or from the sentence with the content that *a* is wrong_2? It strikes me as incredible to claim that one such inference is determinately more rational than another. More plausibly, they are both maximally rational, or it's indeterminate which is uniquely maximally rational.[35] Both options seem relatively plausible. What's important is that on either option 'wrong' vaguely refers.

Importantly, neither explanation of semantic vagueness undercuts our explanation of stability. That's because we explained stability using the impotence strategy, rather than the trumping strategy. So, even if a twin community were, say, strongly disposed to apply their twin term 'wrong' to cases of amputating an innocent bystander's arm to save five people's lives, the referent wrong_1 is not thereby the referent of the twin term. Because their application disposition isn't a factor in fixing reference, their disposition to apply the term to this or that action doesn't "bump up" one referent over another, break the ties between the competing referents, or resolve the indeterminacy as to which referent most rationalizes the practical inferences.

IV. SCHOENFIELD'S REJECTION

Here's where we're at. Schoenfield raises a Challenge: to give a metasemantic theory that predicts that moral terms vaguely refer and that they are stable. If we try to predict vague reference via ties and stability via trumping, we run into trouble. However, CRS gives us an alternative way to predict trumping: the impotence strategy. And it helps us notice an alternative way to predict vague reference: via indeterminate satisfaction. Using the impotence strategy to predict stability, we can predict

35. What if the options are "on par" in Ruth Chang's sense? I think that this case would also lead to indeterminacy in reference. We should revise the reference-fixing criterion of CRS so that terms refer to the referent such that no other referents are more rationalizing than it. If two such relations are on par, then this revised criterion predicts semantic vagueness via ties. See Ruth Chang, "The Possibility of Parity," *Ethics* 112 (2002): 659–88.

vague reference via ties or indeterminate satisfaction without trouble. And because CRS is open to the moral realist, it appears we've rebutted Schoenfield's argument from the Challenge to MVIOV.

Schoenfield, however, thinks that trouble remains. In particular, she thinks that this suggested way of meeting the Challenge is still committed to ontic vagueness. Focusing on the case of CRS predicting vague reference via indeterminate satisfaction, her argument takes two premises:[36]

A1. A robust moral realist who accepts the conceptual role semantics account should think that a [fundamental] language will contain predicates that refer to those properties and relations that make the inference rules for practical reason valid.

A2. If it's indeterminate which properties and relations make the inference rules for practical reason valid, and a [fundamental] language contains predicates that refer to such properties and relations, then a [fundamental] language will contain predicates that lack precise application conditions.

A3. Therefore, the robust moral realist who thinks that it's indeterminate which properties and relations make the inference rules for practical reason valid is committed to ontic vagueness.

Her idea is straightforward enough. A moral realist is committed to there being moral predicates in the fundamental language that we would speak. If we adopt a CRS for moral predicates and CRS leads to vagueness for moral terms, then there is vagueness in the fundamental language that we would speak. Hence, there is ontic vagueness. Note that while Schoenfield focuses on the case of CRS predicting vague reference via indeterminate satisfaction, an analogous argument can be made against the case of CRS predicting vague reference via ties, where 'multiple' replaces 'it's indeterminate which' in the antecedent of A2.[37]

Fortunately, this argument fails for two reasons. First, the argument requires an unreasonable characterization of ontic vagueness. Second, the argument fails to distinguish between moral terms in fundamental languages and moral terms in nonfundamental languages like English. The argument assumes that moral terms in fundamental languages get their meaning via the same conceptual role as moral terms in English.

A. Unreasonable Characterization of Ontic Vagueness

The argument is only valid on Schoenfield's suggested characterization of ontic vagueness:

36. Schoenfield, "Moral Vagueness," 273.

37. Indeed, Schoenfield raises the analogous argument when considering a version of NRM that deploys impotence and vagueness via ties, arguing that we would not be able to make our vague moral terms precise in a fundamental language (ibid., 271).

> **Ontic Vagueness***: There is ontic vagueness just in case, were we to speak a fundamental language, a claim in that language would be indeterminate.

But recall that we gave two reasons that this characterization doesn't "get at" the theoretically interesting phenomenon. Remember, Schoenfield's interpretation of ontic vagueness trivialized the thesis—there are some referentially indeterminate fundamental languages (remember the 'elecuark'/'quarton' example!). The minimum fix we suggested above was to restrict our attention to referentially determinate fundamental languages. If indeterminacy remains in a referentially determinate fundamental language, then the vagueness is genuinely in the world rather than in our language. But on this interpretation of ontic vagueness, the above argument isn't valid, for it doesn't say anything at all about referentially determinate fundamental languages. All the argument shows is that we would speak a referentially indeterminate language about the fundamental.

I can imagine someone responding as follows:

> The 'elecuark'/'quarton' example shows that some fundamental languages are referentially indeterminate. And it's right to say that this isn't enough for ontic vagueness. But Schoenfield's argument, at least in spirit, shows something stronger: that, according to the moral realist, all fundamental languages will have terms that vaguely refer, for all fundamental languages will have moral terms that indeterminately satisfy their associated conceptual roles. Surely this is enough for ontic vagueness. In other words, we should adopt the following sufficient condition for ontic vagueness: if vagueness remains in all fundamental languages, then there is ontic vagueness—even if this is because all fundamental languages are referentially indeterminate!

However, even this repaired argument only goes through if we assume that we are able to speak all fundamental languages—that the notion of fundamental languages relevant for characterizing ontic vagueness is limited to natural languages. But this assumption is implausible: we shouldn't assume that we can speak all fundamental languages—the notion of fundamental languages relevant for characterizing ontic vagueness should include artificial languages that aren't spoken. If we're looking for a description of reality on which no vagueness remains, surely we shouldn't restrict our attention to those languages we can speak. To see this, imagine that there is a referentially determinate fundamental language: the language offers a complete and accurate description of fundamental reality in perfectly precise terms. In addition, no vagueness remains in the language, no claim in the language is indeterminate. But

imagine that, as it happens, that language cannot be spoken by us—it's an artificial language. In what sense is there "vagueness in world"? To focus on natural languages is to take too seriously the imagery of speakers. Indeed, this focus risks trivializing the thesis of ontic vagueness, for it is plausible that all natural languages have vague terms like 'bald' rather than precisely referring terms like 'bald_1' that refer to particular precise properties of the scalp, because no difference in our use could draw fine enough distinctions among these precise properties. (Sure we can gesture at these precise correlates of 'bald', but it strains credulity to think that we could actually speak such a language.)

Once we drop the assumption that all fundamental languages are spoken, it's easy to see where the argument goes wrong. The realist is explaining why moral terms in nonfundamental languages like English are vague yet stable. She does this by appealing to CRS. But CRS is a theory of how the words of natural languages get their reference, not artificial languages. So, even if the reference of moral terms in natural languages is determined by conceptual roles, and conceptual roles induce vagueness, that doesn't show that moral terms in artificial fundamental languages are vague—there are still fundamental languages that give a precise, accurate, and complete picture of fundamental reality, and no claims are indeterminate in such a language.

What is the artificial fundamental language that I have in mind? Recall that the realist thinks that there is a plenitude of precise moral properties $\textit{wrongness}_1$, $\textit{wrongness}_2$, and so on (that was a basic commitment, coming into this discussion, of the realist who denies ontic vagueness). She also thinks that some moral properties are fundamental. If she thinks that the properties in this plenitude are fundamental (the Plenitudinous Fundamentality Option), then there will be an artificial language with a range of moral terms, such as 'wrong_1' and 'wrong_2', and a semantics according to which those terms determinately refer to the precise fundamental properties. Because these precise terms are terms in an artificial language, they aren't associated with any reference-fixing conceptual roles. (And because the language is artificial, it need not be a language we can speak—in the same way that we cannot speak a language with precise terms 'bald_1', 'bald_2', etc., for the scalp.) Alternatively, if she thinks that the plenitudinous properties are nonfundamental and hold in virtue of some precise fundamental property described by some yet-to-be-discovered theory (the Sparse Fundamentality Option), then the artificial fundamental language will have terms and a semantics on which the terms determinately refer to these underlying yet-to-be-discovered properties.

I believe that this is the correct way to respond to Schoenfield's argument: her argument assumes (i) that a referentially indeterminate fundamental language implies ontic vagueness and (ii) that fundamen-

tal languages are natural languages that get a metasemantic treatment like ordinary languages such as English. Neither assumption "gets at" the phenomenon of ontic vagueness. But, to convince those who remain skeptical, I want to point out an additional implausible assumption that her argument rests on.

B. Distinguishing Moral Terms in Fundamental Languages and English

Let's spot Schoenfield that fundamental languages are natural, and that if they are all referentially indeterminate, there is ontic vagueness. Notice that her argument also assumes the following:

> (3) The moral terms in fundamental languages are associated with the same conceptual role as the moral terms in English (or whatever language the moral sorites is presented in).

Cases like Amputations made it overwhelmingly plausible that the English notion of wrongness was vague. If the vagueness was semantic, we needed to show that the English term 'wrong' is referentially indeterminate. We did this by showing that the English term's conceptual role induces vagueness. Of course, if the moral terms in the fundamental language shared this conceptual role, these languages would also be referentially indeterminate. But if the conceptual role for moral terms in fundamental languages is different from the conceptual role for moral terms in English, then indeterminate satisfaction of the conceptual role associated with the English moral term 'wrong' doesn't show that the moral terms in the fundamental language vaguely refer. The role of this assumption is obscured by the generic use of the phrases 'moral predicates' and 'rules of practical reasoning'. Consider Schoenfield's explanation of the cogency of her argument: "So, if conceptual role semanticists are correct in claiming that *moral predicates* refer to the properties and relations that make *the rules of practical reasoning* valid, and it's indeterminate which properties and relations these are, then even in a [fundamental] language, moral predicates will lack precise application conditions. And if *moral predicates* in a [fundamental] language lack precise application conditions, then moral vagueness will exist even among omniscient [fundamental] language users. On such an account, moral vagueness is ontic."[38] However, conceptual role semanticists are merely claiming that *English* moral predicates refer to the properties that make *particular* inferences (in our case, the practical inference to blame) valid (or most rational), and that the terms vaguely refer because of ties or indeterminate satisfaction of the rationalization of these particular inferences. They are not claiming that moral predicates in *fundamental languages* refer to

38. Ibid., 274; emphasis added.

the properties that make these *same* particular inferences valid. In particular, the fundamental moral predicates may be associated with different inferences that do not admit of ties of indeterminate satisfaction.

Indeed, the problem is even worse. For the argument assumes that the correct metasemantics for moral predicates in fundamental languages is given by some version of CRS. But again, we must distinguish treatment of English moral terms from the terms in the fundamental language: CRS may be the correct metasemantics for English while the terms in the fundamental language get their reference through some other means (e.g., NRM).

Given this gap in the argument, there are various ways one might want to fill in this picture. Here are two.

Suppose that the realist accepted Sparse Fundamentality, so she thought that the range of precise moral properties wrong_1, wrong_2, . . . hold in virtue of some yet-to-be-discovered fundamental properties described by some yet-to-be-discovered fundamental moral theory. On this picture, moral fundamentality is sparse: just as a few physical properties give rise to a range of precise nonfundamental properties of the scalp (bald_1, bald_2, . . .), a few fundamental moral properties give rise to a range of precise nonfundamental moral properties (wrong_1, wrong_2, . . .). In this case, the recognition of vague moral terms in English gives us no more reason to think that the terms of the fundamental moral theory are vague than the recognition of the vague term 'bald' gives us for thinking that the terms of the fundamental physical theory are vague. If fundamental moral reality is sparse, then it's easy to come up with metasemantics on which the fundamental moral terms are precise. One way the terms in this fundamental moral theory might be precise assumes that CRS applies to the fundamental language but associates the terms in that language with a different conceptual role than the ones associated with the moral predicates in English. In particular, the yet-to-be-discovered fundamental moral theory will endow the fundamental moral predicates with a theoretical role that (because moral fundamentality is sparse) is determinately played by a unique moral property.[39] Another way posits a different metasemantics for the fundamental language. Suppose, for instance, that NRM applies to the fundamental language. Because the sparse fundamental moral properties are particularly natural, they will exert magnetic pull on the reference of the fundamental moral vocabulary, determinately fixing their reference. (And, of course, because fundamental moral reality is sparse, on either story, small changes in usage of the fundamental term won't shift reference, so the fundamental moral terms are stable.)

39. See David Lewis, "How to Define Theoretical Terms," *Journal of Philosophy* 67 (1970): 427–46.

Suppose instead that the realist accepts Plenitudinous Fundamentality, so that the precise moral properties wrong_1, wrong_2, . . . are all fundamental. So a precise fundamental language will include terms 'wrong_1', 'wrong_2', and so on, that each refer to these precise moral properties. Recall that, according to our CRS for the English term 'wrong', the English term refers to whatever property most rationalizes the practical inference from "*a* is wrong" to the action of blaming those that do *a*, but there were ties or indeterminacy with respect to which precise property had this most rationalizing feature. Of course, if the terms 'wrong_1' and 'wrong_2' are associated with this same conceptual role, they would also be vague. But the obvious response is that these terms will be associated with slightly different practical inferences. According to this explanation, it's most rational for these idealized speakers to treat actions that are wrong_1 ever so slightly differently than actions that are wrong_2. For instance, we might blame someone who performs an action that is wrong_2 ever so slightly more severely than someone who performs an action that is wrong_1, or we might be ever so slightly more wary of such a person, or we might take ever so slightly stronger steps to avoid such actions. If anything like this is the case, then these properties will be associated with ever so slightly different fine-grained "precise" practical inferences. (And because these inferences are practical, the dispositions of these idealized speakers to apply these precise terms are referentially impotent—application dispositions are not a factor that fixes reference. So, stability is maintained.) Indeed, this suggestion fits particularly nicely with the suggestion that the reason the English term 'wrong' is vague is because the notion of blame is vague, thus leading to vagueness in which precise property most rationalizes the practical inference to blame.

One might find the suggestion that we could really speak such a language hard to swallow. That's a reaction I'm sympathetic to. After all, the moral terms play practical roles with vanishingly small differences, differences no real speaker could appreciate. Combined with a skepticism of a sparse, yet-to-be-discovered fundamental moral theory, we might find this whole line of thought unpersuasive.

However, recall that it's a basic commitment of the moral realist who denies ontic vagueness that moral reality contains properties with vanishingly small differences. So, insofar as one thinks that our practices and conventions cannot really draw such vanishingly small distinctions, this seems to illustrate the implausibility of Schoenfield's interpretation of ontic vagueness rather than the plausibility of MVIOV. If our practices and conventions are defective in describing fine-grained reality, this doesn't show that there is ontic vagueness; rather, it shows that natural languages aren't appropriate for characterizing ontic vagueness.

If a skeptic (i) continues to insist that fundamental moral languages must be spoken, (ii) rejects the possibility of sparse, yet-to-be-discovered

fundamental moral theory in favor of a plenitudinous fundamental moral reality, and (iii) thinks that the practices and conventions of idealized speakers cannot really draw such fine-grained distinctions, then the skeptic should accept MVIOV. But this acceptance of MVIOV no longer has much to do with the Challenge—it is not derived from reflecting on the stability of English moral terms. Thus, we'll still have shown that the stability of moral terms creates no special problem for a realist who thinks that moral vagueness is semantic.

V. CONCLUSION

Our Challenge was to find a metasemantic theory that predicted that it is vague what English moral terms refer to and yet preserved stability. So long as we thought that a metasemantic theory predicts semantic vagueness via ties in the reference-fixing criteria and predicts stability via the trumping strategy, it was hard to see how we could meet that Challenge. But reflecting on CRS-based metasemantic theories led us to alternative strategies for predicting stability (the impotence strategy) and semantic vagueness (indeterminate satisfaction), providing us with a way to meet that Challenge. Schoenfield rejected this path to meeting her Challenge, but her rejection depended on an implausible characterization of ontic vagueness and a failure to distinguish between the reference-fixing criteria of moral terms in English and moral terms in fundamental languages.