

Moral Vagueness is Ontic Vagueness

* Many thanks to Elizabeth Barnes, Ross Cameron, Jamie Dreier, Dan Korman, and especially to Matti Eklund, Tom Dougherty and Susanna Rinard for helpful discussion and for reading and commenting on earlier drafts. I also received extremely helpful and detailed feedback on this paper from two anonymous referees, and from audiences at the 2014 meeting of the *American Philosophical Association* – Pacific Division, the faculty reading group at The University of Texas at Austin, and the “Indeterminacy in What We Care About” conference held at the University of Leeds.

Abstract

The aim of this paper is to argue that, if a robust form of moral realism is true, then moral vagueness is ontic vagueness. The argument is by elimination: I show that neither semantic nor epistemic approaches to moral vagueness are satisfactory.

1. Introduction

James Hutton Kidd, in 1924, developed his own variety of apple by crossing Red Delicious apples with Cox’s Orange Pippin variety. The resulting apple is sweeter than the Cox variety, more chewy than crunchy, and red. Well, maybe not red. Orange perhaps? Kidd himself couldn’t make up his mind and the apple was named “Kidd’s Orange Red.” Suppose I give you such an apple. I let you touch it, look at it, and measure its reflectance properties. None of this investigation, however, settles for you the question of whether the apple is red. But note that, even if you are a great apple connoisseur, the question of whether Kidd’s apples are red will probably not keep you up at night.

Sometimes, however, we *are* very interested in knowing whether or not something is red. Was the traffic light on the corner red (upon having driven through the intersection and noticed a camera mounted on top)? Is the flag waving on the beach red (while trying to determine whether the water conditions are too dangerous for swimming)? Is the belly of the Eurasian bullfinch red (while trying to identify the species of a bird you just saw)?

Plausibly, the reason that these redness questions are so much more interesting than the color of Kidd's apples is that your puzzlement about the color of Kidd's apples arises from the vagueness of the predicate "red," while this is not so in the other cases described. Indeed, many philosophical accounts of vagueness nicely explain why you wouldn't care about the color of Kidd's apples. Popular semantic accounts of vagueness tell us that we just haven't gotten around to settling on a precise cutoff point for "red." Many candidate cutoff points would be acceptable and, if only we spoke more precisely, our word "red" would refer to a precise property. Some epistemicist accounts tell us that, despite our lazy linguistic behavior, there is a precise cutoff point for "red." We just can't know where the cutoff point is because the location of the cutoff is extremely sensitive to linguistic usage. A few more people calling Kidd's apples "red" results in the referent of "red" changing. Both of these accounts nicely explain why your apple enthusiasm doesn't lead to a deep curiosity about the color of Kidd's apples. For, given everything you already know, the sorts of considerations that might settle the question of Kidd's apples for you, if it could be settled at all, are merely linguistic.

However, we're not always so nonchalant about puzzlement due to vagueness. It is plausible that moral predicates are vague (more on that later), and we certainly care deeply about whether, for example, in some potentially borderline case, an act is *permissible*.

Unlike the color of Kidd's apples, whether such an act is permissible *is* the sort of question that will keep us up at night. In this paper I will argue that many of the accounts of vagueness that serve us well when we focus on properties such as redness and baldness aren't well suited to account for vagueness in what matters, in particular, moral vagueness. More specifically, I will argue that if moral realism is true, then the only satisfactory account of moral vagueness is an ontic account: roughly, an account according to which the vagueness stems from *how things are in the world* (and not how we describe the world, or what we know about it).

This thesis is important for three reasons. First, moral theorists need an account of moral vagueness since Sorites-like reasoning is used in a variety of arguments for first order moral claims. Consider, for example, the following worry for deontological views described by Anthony Ellis¹ and Larry Alexander²: if the deontologist is going to avoid the unattractive consequence that there is no number of lives for which it would be permissible to trade off a single person's mild discomfort, they must think that there is a *threshold* at which the consequentialist considerations become overriding. But, they claim, such a threshold would be unacceptably arbitrary. Another example of Sorites-like reasoning appears in the debate between Alastair Norcross³ and Larry Temkin⁴ concerning whether the betterness relation is transitive. Temkin, in defending the *intransitivity* of betterness, claims that there can be a series of lives: A, B...X, Y, in which A is better than B, B is better than C, ... X is better than Y, but in which Y isn't better than A – in fact, it is *worse* than A. Norcross, in arguing against this possibility, claims that Temkin is committed to the existence of some property that *explains* what makes A worse than Y. And, in the course of his argument, Norcross relies on the claim that there will be “a *first life* in the sequence that

does not contain the property” (my emphasis).⁵ The success of these arguments depends on the right account of moral vagueness: if moral vagueness is, for example, ontic, then it’s not clear that the relevant *threshold* or *first life* will exist.

Second, the claim that ontic vagueness exists is highly controversial. If, as I will argue, moral realism is committed to the existence of ontic vagueness, this will give opponents of ontic vagueness a reason to reject moral realism. Thus, thinking about how to properly account for moral vagueness could provide new traction on central debates in metaethics.

Finally, for those (like myself) who are more open minded about ontic vagueness, this paper shows that if we’re looking for a fully general account of vagueness, we need to consider a broader range of cases than those typically discussed. Vagueness can pop up in much more interesting places than balding men’s heads, and some of these more interesting forms of vagueness cannot be well accounted for in the usual ways.

2. The Central Thesis

The central thesis of this paper is the following:

Central Thesis: If a robust form of moral realism is true, and there is moral vagueness, then it is ontic vagueness.

Let me begin by clarifying some of the terminology. First, when I talk about moral vagueness I will be talking primarily about vagueness for moral *permissibility*.

Second, by “a robust form of moral realism” I mean a view according to which moral

truths are necessary and, following Tom Dougherty, according to which moral properties “are part of the deep underlying metaphysical structure of the world, and they obtain entirely independently of how we conceptualize the world.”⁶

Finally, I will follow Elizabeth Barnes’ characterization of ontic vagueness as:

“[Vagueness] in how things are... [vagueness] that would remain even if we spoke a perfect language and were omniscient.”⁷

By “a perfect language”, Barnes means a language that contains all and only predicates that are necessary to provide a complete and accurate description of how things are fundamentally.

Barnes’ criterion is certainly not the only way of thinking about ontic vagueness and my goal here is not to adjudicate between different ways of using the term “ontic vagueness.” I use her criterion simply because it gets at the phenomenon I’m interested in. What I want to know is the extent to which *morality itself* is responsible for moral vagueness. If a complete and accurate description of how things are fundamentally contains vague predicates, it is plausible that this is *because* of how things are fundamentally. For note that one can imagine a way that things could be such that the best description of how things are fundamentally would contain only precise predicates. So if, as a matter of fact, the best description of how things are fundamentally contains vague language, our world must not be of that sort. In other words, if our world, at the fundamental level, is best described *imprecisely*, it is at least partially the world’s fault that things are so. When I say that moral vagueness is ontic I’m not committing myself to any particular view of ontic vagueness (vague objects, vague properties, etc.). All I’m claiming

is that moral vagueness exists because of how things are *morally*, independently of how we describe or conceptualize the world, and independently of what we know about it.

The argument for my central thesis will be an argument by elimination. I will argue that if moral realism is true, moral vagueness cannot be well accounted for by purely semantic or epistemic theories of vagueness. There are two challenges that any such argument will face: the first is that there might be views of vagueness that don't fit neatly into the semantic/epistemic/ontic trichotomy. Second, arguments by elimination frequently remain conveniently silent on the uneliminated option, and mine is no exception. There are well known challenges to ontic vagueness,⁸ which I will not be addressing. So one may well worry that showing that semantic and epistemic accounts fail isn't enough to show that an ontic account succeeds.

My response to these worries is quite concessive. Regarding the first worry: I will provide no guarantee that *no* non-ontic account of vagueness is possible. Rather, I will show that many accounts of vagueness in the literature, which *can* be described as either semantic or epistemic, do not extend well to the moral case, specifically because, if moral realism is true, moral predicates lack either the semantic or epistemic features that are thought to be essential to vague predicates. However, it is worth noting that given the definition of ontic vagueness being employed in this paper, it's natural to think that something like the familiar semantic/epistemic/ontic trichotomy *does* obtain. Recall that I follow Barnes in defining ontic vagueness as vagueness that would remain even if *we spoke a perfect language and were omniscient*. If one denies that moral vagueness is ontic, then one thinks that moral vagueness is *not* a phenomenon that would exist if we spoke a perfect language and were omniscient. Now, if moral vagueness exists amongst us, but

wouldn't exist amongst omniscient perfect language users, this is presumably because we either fail to speak a perfect language, or are not omniscient. Thus, we can divide non-ontic accounts into those according to which it is in virtue of our *language* that vagueness exists, and those according to which it is in virtue of our *ignorance* that vagueness exists.

Still, I cannot claim to have covered every inch of logical space. After all, one might have a view according to which it is a *combination* of our linguistic and epistemic situations that makes for vagueness. Additionally, for both the semantic and epistemic categories, I will focus on the sorts of accounts that have been prevalent in the literature. But even if I show that all such accounts fail in the moral case, one might think that there is some *other* yet to be identified feature of our language or our epistemic situation that constitutes moral vagueness but that wouldn't exist in a community of omniscient perfect language users. So I encourage readers who are interested in developing a non-ontic account of moral vagueness to think about the degree to which the features of vague predicates relied on in the semantic and epistemic accounts prevalent in the literature are essential to a non-ontic theory of vagueness.

Regarding the second worry, I will, indeed, say nothing *positive* in this paper in defense of ontic vagueness. I can point the reader to some excellent literature in its defense,⁹ though there is also, of course, excellent literature opposing ontic vagueness.¹⁰ So I do not hope to convince the unconvinced. Those who are troubled by ontic vagueness can instead think of my argument by elimination as a trilemma for moral realism,¹¹ or for the existence of moral vagueness. For if robust moral realism has the consequence that semantic, epistemic *and* ontic accounts of moral vagueness are inadequate, then, if the traditional tripartite division is exhaustive, either robust moral realism, or the claim that

there is moral vagueness must be rejected. So there are different lessons one might draw from the arguments I will present. Since I do not find ontic vagueness so troublesome, and I find robust moral realism plausible, I will present the arguments as favoring an ontic account of moral vagueness. But the considerations I raise may lead other theorists down other equally interesting paths.

3. Moral Vagueness

I will not give a full defense of the claim that moral vagueness exists, since I am only interested in defending the conditional claim: that *if* moral vagueness exists, and robust moral realism is true, the vagueness is ontic.¹² I will, however, provide some cases that illustrate the phenomenon that I am interested in.

DIVERSIONS: Darryl is watching his two year old daughter play in a city park. It is permissible to divert his attention from her for 1 second. It is not permissible to divert his attention from her for 5 minutes. Is it permissible to divert his attention for 30 seconds? 31? 32? Plausibly, we can create a Sorites series, admitting of borderline cases of permissibility, out of a series of diversions whose lengths differ by a second.¹³

ABORTIONS: Cheryl is pregnant. She and her partner suddenly realize that, if the pregnancy is carried to term, they'll have to skip a much-anticipated vacation that they had long ago planned with a group of friends. They don't think that skipping the vacation would have a significantly negative impact on their lives. But, all things

considered, they'd prefer that the fetus not be born. It is permissible to abort after one day for these reasons. It is not permissible to abort after nine months for these reasons. Is it permissible to abort at 150 days? 151? 151.5? Plausibly, we can create a Sorites series, admitting of borderline cases of permissibility, out of a series of abortions in which the fetus's age differs by a day (or a minute, or a second).

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.

The particular moral judgments in any given case may be controversial. However, I suspect that, on most plausible moral theories, there will be at least some vagueness concerning what is permissible.¹⁴

4. The Argument for the Central Thesis

In this section I will argue that both semantic and epistemic treatments of moral vagueness are problematic for the robust moral realist. I will conclude that the best account of moral vagueness for the moral realist is ontic.

4.1. Why Semantic Treatments of Moral Vagueness are Problematic

According to a *semantic* conception of moral vagueness, moral vagueness is a feature of our language. But what feature? On common semantic conceptions, vague predicates are predicates that lack precise application conditions. Because the reference-fixing facts fail to determine which, of a myriad of precise properties, a vague predicate refers to, some sentences containing vague predicates are neither true nor false. If moral vagueness is not ontic, it follows that these imprecise application conditions exist only in imperfect languages. Should the moral realist accept a semantic account?

The first thing to note is that the moral realist is committed to the thought that a perfect language *will* contain moral predicates (since such predicates, for the realist, are necessary to give a complete and accurate account of how things are fundamentally). So the moral realist who thinks of moral vagueness as a purely semantic phenomenon thinks that in a perfect language moral predicates would exist, but they would have precise application conditions. I will argue that this view is problematic. My argument against the semantic conception of moral vagueness takes the form of a dilemma. There are two versions of semantic accounts that I will consider, which I will call “shifty” and “rigid,” respectively. I will argue that both of these views face serious difficulties.

The distinction between shifty and rigid views is best brought out by considering a community very much like ours, called “Liberal.” The only difference between our community and Liberal is that the liberals use the word “permissible” slightly more liberally than we do. In particular, they unhesitatingly apply the term “permissible” to 50-second diversions in cases like Darryl’s, whereas we (suppose) unhesitatingly apply the term “impermissible” to 50-second diversions. (Fill in the details of the case however you like to get the verdict that diverting for 50 seconds is impermissible. Whether or not a

given number of seconds will constitute an impermissible diversion will depend on many factors including the number of other adults in the park, the maturity of the child, the time of day, and so on). Now consider the following question: Is the sentence: “A 50 second diversion is permissible in a case like Darryl’s” true when uttered by a speaker of Liberal? The shifty view says yes. The rigid view says no.

The Shifty View: The truth-value of an utterance: “X is permissible” is highly sensitive to the way the word “permissible” is used in a linguistic community. A sentence “X is permissible” may have one truth-value when uttered by S, in English, but a different truth-value when uttered by S’, who is in all respects just like S, except that S’ lives in a linguistic community that applies the predicate “permissible” slightly more liberally than we do.

The Rigid View: The shifty view is false. Slight changes in the application of “permissible” will not impact the truth-values of sentences containing this predicate.

The Shifty View

Let’s call a member of Liberal, who, like Darryl, is watching his daughter play at a similarly situated park, Darryl*. What could explain why, when Darryl says: “it is permissible to divert my attention for 50 seconds,” he speaks falsely, but Darryl*, uttering that same sentence, speaks truly? The moral realist cannot say that what explains this is the fact that Darryl* is permitted to divert his attention for 50 seconds, while Darryl is not, and that this difference in permissibility is a result of the different practices in their

respective communities. For the moral realist, an act doesn't become permissible just because one's community sanctions it.¹⁵ So if Darryl*'s utterance is to be true, it must be because his utterance expresses a different proposition than Darryl's. This will be explained by the shiftiness of the referent of "permissible." Due to the linguistic behavior of the Liberals, their *word* "permissible" has a slightly different extension from the one picked out by our word "permissible." This view is perfectly consistent with (at least the letter of) moral realism. It maintains that, despite the practices in Darryl*'s community, Darryl*'s actions are impermissible. Nonetheless, Darryl*'s *utterance*: "diverting for 50 seconds is permissible" is true.

The problem with the shifty view is that, at least for a moral realist, it can't make good sense of moral deliberation. Suppose that Cheryl and her partner are deliberating about whether to abort a fetus at 150 days. They feel very conflicted about the issue and they spend a great deal of time deliberating, indeed, *agonizing*, over whether such an abortion would be permissible. The linguistic anthropologist then knocks on the door. "Guess what!" she says. "I've conducted a series of surveys about the way language users in your community use the word "permissible." Here is the data!" After dropping some thick manila folders on the coffee table, the anthropologist disappears. Fortunately, Cheryl and her partner are expert philosophers of language and they can make excellent inferences about the truth-values of sentences with vague predicates based on usage facts. Cheryl and her partner spend the night crunching through the data that the linguistic anthropologist provided. With the first rays of light, Cheryl and her partner breathe a sigh of relief. The usage facts in their community are only consistent with precisifications that permit the abortion in question. Thus, the abortion is permissible.¹⁶

Note that the claim that Cheryl can learn what is permissible by crunching through the data doesn't mean that what is permissible *depends* on linguistic usage, in the sense that, had we used language differently, different things would be permissible. What does, however, follow from the shifty semantic account is that Cheryl can find out that some abortion whose permissibility she was uncertain about, is, in fact, (determinately!) permissible by collecting linguistic data.¹⁷ However, it doesn't seem like crunching through linguistic data is a way of resolving doubts about the permissibility of abortion, especially for the moral realist. Linguistic anthropologists may be helpful with all sorts of things, but solving moral conundrums is not one of them.

Note that these considerations don't generalize to other predicates for which a shifty semantic treatment seems plausible. Consider, for example, the word "bald." Unlike permissibility, nobody agonizes about borderline cases of baldness, and if someone did bother to deliberate over such a silly question, it seems perfectly sensible that they could conclude their deliberation by gaining usage information from the linguistic anthropologist. But the way we deliberate about morality is different from the way we deliberate about baldness. And a shifty semantic account of moral vagueness can't make good sense of the fact that our moral deliberations in difficult cases cannot be resolved by learning usage facts.¹⁸

A quick note before proceeding: even if it seems odd for a realist to think that the linguistic anthropologist can help us resolve moral deliberation, one might wonder whether this puzzling consequence is a special consequence of the shifty semantic view of moral vagueness. After all, as long as we accept that *x* is permissible if and only if "*x* is permissible" is true, then learning the referent of "permissible" can resolve our

deliberation about what is permissible, regardless of how we account for moral vagueness. Isn't this a problem for everyone?¹⁹ The answer is no. This is because, as I will illustrate below, on non-shifty views, we cannot resolve our moral deliberation by learning the referent of "permissible" through usage facts alone. We'd *also* have to learn substantive moral facts. What's special about the shifty view is that we can gain information that is relevant to our deliberation *just* by learning facts about linguistic usage.

To see why non-shifty views don't face the linguistic anthropologist problem, consider, for example, Ralph Wedgwood's conceptual role semantics for moral predicates.²⁰ On this view, "permissible" will refer to whatever property is best suited to play the characteristic role in practical reasoning. On Wedgwood's account, minor variations in linguistic usage aren't going to affect the referent of "permissible" because minor variations in usage won't change what property is best suited to play the characteristic role. On this non-shifty view, Cheryl couldn't resolve her deliberation by gaining data from the linguistic anthropologist. She'd need the *philosopher* to come knocking at the door to tell her which property was best suited to play the characteristic role in practical reasoning. And there is nothing wrong with thinking that *this* sort of information can resolve moral deliberation.

However, even those who reject conceptual role semantics, and favor, say, a causal semantic theory, could avoid the problematic consequence if they didn't need to appeal to shiftiness to account for moral vagueness. If one thought, as our robust moral realist does, that moral properties are "part of the deep underlying metaphysical structure of our world,"²¹ then moral properties could act as reference magnets. According to the doctrine of reference magnetism, certain properties are better candidates for the reference of our

terms than others. You might have two different semantic theories, T1 and T2, that fit equally with a community's usage and yet T1 is correct while T2 is not. This would be the case if T1 has our terms refer to properties that are more eligible for reference (hence "reference magnetic") than T2. On David Lewis's view,²² natural or "joint-carving" properties are especially well suited to act as reference magnets. A robust realist who subscribed to reference magnetism could adopt the view that any linguistic community with usage sufficiently similar to our own would use the word "permissible" to refer to the same property that we refer to. The "magnetic pull" of this property would attract the reference of the word across a variety of usage patterns. Although it is true, on this picture, that learning the referent of "permissible" would tell us what is permissible, this would involve learning which of the potential referents of "permissible" is the metaphysically privileged one, the one that carves nature at its normative joints. The thought that this sort of information could help resolve deliberation about permissibility is not so puzzling.

In sum, what's odd is not the thought that learning the referent of "permissible" could, in certain circumstances, resolve moral deliberation. What's odd is thinking that learning detailed information about *linguistic usage* could resolve moral deliberation. This consequence only arises for views according to which the referent of "permissible" shifts with minor variations in usage.²³

I will end by briefly noting that a shifty semantic view may also face problems along the lines of those described by Horgan and Timmons.²⁴ Cases of apparent disagreement about permissibility between us and the members of Liberal will turn out to be cases in which we are talking past each other, since the referent of our term shifts with minor

variations in usage. Matti Eklund²⁵ raises worries related to these for shifty semantic views of moral predicates.

The Rigid View

In the previous subsection I argued that the moral realist shouldn't accept a shifty semantic account because such an account has the bizarre consequence that learning detailed information about usage could resolve moral deliberation. I will now argue that rigid semantic views should be rejected as well. Since every semantic view is shifty or rigid, this will conclude my argument against semantic accounts of moral vagueness.

According to the rigid view, even if everyone in Darryl*'s community applies "permissible" to 50-second diversions, Darryl*'s utterance: "diverting for 50 seconds is permissible" is false. To evaluate the plausibility of the rigid view we need to think about why, in general, changes in linguistic usage might not result in a different referent. After all, most vague predicates *are* shifty. If people in Liberal used the word "red" or "bald" more liberally, the words "red" and "bald" as uttered by them, would *not* have the same extension as our words "red" and "bald."²⁶ So if, in contrast with other vague predicates, "permissible" isn't shifty, we need some explanation of why this is so. In discussing the shifty view, I mentioned two potential explanations. A conceptual role semantic theory might induce rigidity and a theory in which the moral properties were reference magnetic might induce rigidity. Neither of these views, however, can successfully explain the rigidity of vague moral predicates.

Why is this? If *one* precisification of "permissible" were most natural, or best suited to play the practical reasoning role, and this explained why shifts in usage don't result in a

shift in referent, this would be because “permissible” referred to that one precisification across a variety of usage patterns. But if “permissible” refers to just one precisification, then “permissible” has precise application conditions, contrary to the claim of the semantic theorist. So if there is semantic indeterminacy, the rigidity of moral predicates can’t be explained by appealing to *one* property as being the most natural or best suited to play the practical reasoning role. The story will have to be more complicated. Below I will consider how both reference magnetism and conceptual role semantics might account for rigidity in the presence of semantic vagueness.

Reference Magnetism as an Explanation for Rigidity

Dougherty²⁷ describes a view according to which there is a *set* of equally natural precise properties that are candidate referents for “permissible.” Suppose a rigid semantic theorist adopted this view. She could then think that naturalness considerations constrain what properties can be acceptable precisifications of “permissible”. If P is not sufficiently natural, then, even given certain changes in linguistic usage, P will not be a property such that it is indeterminate whether “permissible” refers to P.

Dougherty’s suggestion raises the following question: might the linguistic usage of a community determine which members, *amongst those that are tied for naturalness*, a community is semantically undecided between? (We’ll call the set of properties a community is undecided between “the range of indeterminacy.”)

If the answer is yes, the view has turned into a shifty view. The reason is as follows: Suppose that $P_{30} \dots P_{40}$ are precise properties that are sufficiently natural, and that the subscripts correspond to the number of seconds that will constitute the cutoff point for

permissible diversions in cases like Darryl's. If the linguistic usage of a community can narrow down the range of indeterminacy, then there might be two communities, A and B, with the following features:

- In virtue of Community A's linguistic usage, the word "permissible" in Community A determinately *excludes* from its extension diversions of 35 - 40 seconds, and it is indeterminate whether diversions of 30 - 35 seconds are in the extension of "permissible."
- In virtue of Community B's linguistic usage, the word "permissible" in Community B determinately *includes* in its extension diversions of 30 - 35 seconds, and it is indeterminate whether diversions of 35-40 seconds are in the extension of "permissible."

Any view which deems such cases possible is a shifty view: one according to which the extension of "permissible" shifts due to linguistic usage. But now that we're back in shifty territory, we will face the problems that shifty views face. For example, if you were unsure whether your linguistic community had usage patterns that were more like Community A or more like Community B, you could find out that some particular diversion is permissible by gaining data from the linguistic anthropologist.

So let us suppose that the answer is no. There is a set of equally natural precise properties that compose the range of indeterminacy for "permissible" and narrowing the class of things that we call "permissible" won't narrow the range of indeterminacy. In other words, even if our usage consistently excluded, say, 40-second diversions, our word "permissible" would still indeterminately refer to properties that have 40-second diversions in their extension.

The first thing to note about such a view is that it is committed to ontic vagueness. If narrowing the class of things we call “permissible” cannot narrow the range of properties that we indeterminately refer to, it’s hard to see how moral predicates, even as used by speakers of a perfect language, could fail to be vague. For on the view in question, even if there were universal and unhesitating agreement amongst members of the community about every case, the predicate would *still* indeterminately refer to all of the properties that we indeterminately refer to. But if universal unhesitating agreement isn’t going to make a semantically vague predicate precise, it’s unclear what features of language use *could* make our moral language precise.

The second thing to note about the rigid view is that it is not delivered by any plausible theory of reference magnetism. For reference magnetism is generally thought to be a phenomenon that directs our reference from less natural to more natural properties – not a phenomenon that forces us to refer, or indeterminately refer, to *every* equally natural property in the vicinity of the properties that are consistent with our usage. At least on common understandings of reference magnetism, if the usage of all of the predicates in our language coincided perfectly with a set of perfectly natural properties, reference magnetism would have no work to do. It wouldn’t make us indeterminately refer to other perfectly natural properties that are not consistent with our usage. So a view according to which “permissible” lacks precise application conditions, but usage can’t narrow down the range of indeterminacy due to reference magnetism, is not a viable one. At very least, one would need a non-traditional story about how reference magnetism worked.

Conceptual Role Semantics as an Explanation for Rigidity

Let's now consider whether the rigid semantic view could be explained by appealing to conceptual role semantics. For ease of discussion, I am going to shift the focus from the predicate "permissible" to the predicate "all things considered better than" simply because the latter is the predicate that Ralph Wedgwood,²⁸ who develops such a semantics in detail, uses. Wedgwood thinks of the relation named by this predicate as a four-place relation: x is all things considered a better thing than y for z to do at t . Plausibly, this relation is vague. For recall the Sorites series generated by AMPUTATIONS. One might reasonably think that, taking into account all the relevant factors, *it's all things considered better* for a subject to amputate an arm in order to save a billion lives, than it is to avoid such an amputation and let the billion die. And one might also reasonably think that it's *not* all things considered better for a subject to amputate an arm to save *one* life, than it is to avoid such an amputation and let the one die. As with permissibility, we can construct a Sorites series of cases in which you are forced to choose between amputating an arm and saving an increasing number of lives. Such a series will admit of borderline cases, just like the Sorites series using "permissibility" does.

On Wedgwood's view, "better than" will refer to whatever relation makes certain rules of inference valid. In particular, he thinks that the referent of "better than" will be that relation which makes the inference from accepting the sentence: "It's better for me to do x than y at t ," to forming the preference for doing x over y at t , a valid one. Note that these rules of inference are rules of inference for *practical* reason since they sometimes output a preference, rather than a belief. What does it mean to say that an inference outputting a preference is valid? The notion of validity that Wedgwood has in mind is the notion of *correctness* preservation, where having a correct preference, according to

Wedgwood, is having a preference that is permitted by the goals of practical reason. Now, 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. Let's consider each possibility in turn.

On Wedgwood's view, there is no indeterminacy in which inference rules constitute the conceptual role for "better than." The referent of "better than," for Wedgwood, is (determinately!) "that four-place relation between x , y , z , and t , such that, necessarily, it is correct for z to prefer doing x over doing y at t and a mistake for z to prefer doing y over doing x at t if, and only if, x , y , z , and t stand in that relation."²⁹ But perhaps Wedgwood is mistaken. For instance, one might think that it is indeterminate whether "better than" stands for a *subjective* notion of betterness (one that takes our ignorance about descriptive matters into account), or an *objective* one (one that doesn't). But even if it is indeterminate which rules of inference are constitutive of the conceptual role for "better than," such indeterminacy will not explain the indeterminacy that arises from the Sorites series in cases like AMPUTATIONS. For if the vagueness that arises in AMPUTATIONS were a result of indeterminacy about which rules constituted the conceptual role for "better than," then fixing what the rules are would fix which precise relation is the referent of "better than." But this is implausible. Fixing the rules of practical reasoning that define the conceptual role for "better than" (say we use the subjective notion) *doesn't* determine a cutoff in the Sorites series for AMPUTATIONS.

A more promising solution, then, is to propose that the indeterminacy in the referent of moral predicates is a result of the fact that *it's indeterminate which relations*

make the inference rules for practical reason valid. But if it is indeterminate which relation makes the inference rules for practical reason valid, there is ontic vagueness, and so the account is not purely semantic. I will demonstrate this by arguing for the following two claims:

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

(B) If it is indeterminate which properties and relations make the inference rules for practical reason valid, and a perfect language contains predicates that refer to such properties and relations, then a perfect language will contain predicates that lack precise application conditions.

From (A) and (B) it follows that the robust moral realist who thinks that it's indeterminate which properties and relations makes the inference rules for practical reason valid is committed to ontic vagueness. (For recall the Barnes test: Ontic vagueness is vagueness that would remain even if we spoke a perfect language and were omniscient. If a perfect language contains predicates that lack precise application conditions, vagueness would remain even in the presence of omniscient speakers of this language).

Why is (A) true? As I mentioned earlier, the robust realist must think that a perfect language contains predicates that refer to moral properties and relations. For recall that a perfect language is a language that contains all and only predicates that are necessary to give a complete and accurate description of how things are fundamentally. If a perfect language contained no moral predicates, then a complete and accurate description of how things are fundamentally wouldn't tell us how things are morally. Now, since, according to

conceptual role semantics, the moral properties and relations are identical to the properties and relations that make certain rules of practical reason (call them *R*) valid, it follows from the fact that a perfect language will contain moral predicates that a perfect language will contain predicates that refer to those properties and relations that make *R* valid.

Why is (B) true? Suppose that a language contains predicates that refer to the properties and relations that make *R* valid, and that it is indeterminate which properties and relations these are. Then, such predicates will lack precise application conditions. For if the predicates referred to properties and relations with precise application conditions, it would follow from conceptual role semantics that *that these precise properties and relations* make the rules of practical reason valid. But this is inconsistent with the claim that it's indeterminate which properties and relations make the rules of practical reason valid.

So, if conceptual role semanticists are correct in claiming that moral predicates refer to the properties and relations that make the rules of practical reason valid, and it is indeterminate which properties and relations these are, then, even in a perfect language, moral predicates will lack precise application conditions.³⁰ And, if moral predicates in a perfect language lack precise application conditions, then moral vagueness will exist even amongst omniscient perfect language users. On such an account, moral vagueness is ontic.

Let's retrace the dialectic. We were considering semantic accounts of moral vagueness. I argued that *shifty* semantic accounts – those according to which the referent of “permissible” shifts with minor variations in usage – should be rejected. Thus, if a semantic account of moral vagueness is to succeed it must be a *rigid* semantic account: one according to which, even if a community consistently applies the predicate “permissible” to

impermissible acts, the referent of the term doesn't shift. I noted that a rigid semantic account of vagueness is unusual. Most vague predicates get a shifty semantic treatment. So what's special about moral predicates? The explanation considered in this subsection is that the referents for moral predicates aren't determined by which acts we apply the predicate to. Rather, they are determined by which properties make the rules of practical reason valid. (This is in contrast to predicates like "red" and "bald" whose referents *are* determined primarily by what sorts of things we apply the predicate to). So the suggestion on the table was this: since applying "permissible" to a different class of instances doesn't change what makes the rules of practical reason valid, shifts in usage won't result in shifts in reference. Hence, the rigid view. This all works splendidly if there is no indeterminacy. However, I argued, if there is moral vagueness, the conceptual role theorists must think that it is indeterminate which properties and relations make the rules of practical reason valid. I then argued that any account according to which it is indeterminate which properties and relations make the rules of practical reason valid is committed to ontic vagueness.

The discussion above demonstrates that an account of moral vagueness can be both semantic *and* ontic. (Indeed, as Barnes notes: "[m]ost any plausible example of ontic vagueness would likely be an example of *both* semantic and ontic vagueness. The claim of ontic vagueness is thus best understood as this: the *source* of at least *some* of the vagueness is ontic, not semantic").³¹ My goal, remember, is to argue that moral vagueness is ontic vagueness. If it turns out that there are semantic views that are also ontic, the correctness of such a view is perfectly consistent with my thesis. Not only is the conceptual role view under consideration both semantic and ontic, it is a view according to which there is

semantic indeterminacy *because* of ontic indeterminacy. It is because the world doesn't determine which properties make the rules of practical reason valid, that it's indeterminate which properties our moral predicates refer to. In light of these considerations, I need to amend the first premise of the argument for my central thesis. I needn't commit myself to the view that semantic accounts of vagueness are problematic for the moral realist. I only need to argue that what we might call "purely" semantic accounts of vagueness – accounts according to which there is semantic vagueness but no ontic vagueness – are problematic for the moral realist.

In sum, rigid semantic views require explanation. Most semantically vague predicates are shifty, so why would moral predicates be rigid? Reference magnetism and conceptual role semantics can explain rigidity if there is no indeterminacy. But if there is indeterminacy, then, even if these theories could yield the required kind of rigidity, the resulting view will be a view on which there is not *only* semantic vagueness, but also ontic vagueness. The shifty semantic view, on the other hand, can't make good sense of moral deliberation. Since any semantic view must be shifty or rigid, I conclude that a (non-ontic) semantic view cannot provide a satisfactory account of moral vagueness.

4.2. Why Epistemic Treatments of Moral Vagueness are Problematic

I will say that an account of vagueness is *epistemic* if it is an account according to which vagueness is inherently an epistemic phenomenon and borderline cases are characterized as cases in which a special sort of ignorance is involved. This way of understanding epistemic accounts insures that such an account is not ontic. Since

vagueness is essentially ignorance, in a perfect linguistic community of omniscient agents there will be no vagueness, since there will be no ignorance.

According to epistemicist accounts of vagueness (such as Timothy Williamson's)³² vague predicates have precise application conditions but it's extremely difficult, perhaps even impossible for agents like us, to know what they are.³³ For the purposes of this section, I want to set aside the worries one might have about the plausibility of epistemicism in general. Rather, I want to focus on specific worries that arise with extending the epistemicist view of vagueness to *moral* predicates. My concern about epistemicism as an account of moral vagueness will be that the *explanations* that epistemicists have offered for why it's so difficult to know the location of the precise cutoff points are not available to the moral realist. Let me begin with Williamson's explanation:

Although someone may judge truly 'Baldness is the property of having fewer than 3,832 hairs on one's scalp', the judgment does not express knowledge, for whatever produced a judgment in those words could very easily have done so even if the overall use of 'bald' had had been very slightly shifted (as it could very easily have been) in such a way that it referred to the property of having fewer than 3,831 hairs on one's scalp, in which case the judgment then made in those words would have been false.³⁴

Call this "the safety-based explanation."³⁵ It is a presupposition of this explanation that, had the overall use of vague predicates been "very slightly shifted", they would have had different referents. But recall that views of moral predicates that have this shifty feature cannot make good sense of moral deliberation. The Williamsonian explanation yields the result that Cheryl could (in principle, though it would be extremely difficult!) resolve her

deliberation about whether aborting her fetus is permissible by learning enough about her community's linguistic usage. But this is implausible. Thus, Williamson's explanation for the unknowability of the sharp cutoff point is not available to the moral realist. Although the view is an epistemic view, the explanation for the epistemic phenomenon relies on a shifty semantics, which, I have argued, is problematic for moral realism.

At this point one might wonder how central Williamson's safety-based explanation is to the epistemicist's project overall. Roy Sorensen, for example, at one point suggested that it is not necessary to provide any explanation for the fact we cannot know the location of the sharp cutoffs.³⁶ If the explanation given by Williamson for the unknowability of these cutoffs is just the icing on the epistemicist's cake, then perhaps the moral epistemicist should simply excise this feature of Williamson's picture and hold on to the rest. In fact, however, the explanation plays a very important role. Stephen Kearns and Ofra Magidor provide a nice explanation for why this is so:

It is important to note that [the explanation of why the sharp cut-off points for vague predicates are unknowable] plays a central role in Williamson's defense of the epistemicist position, for it provides a response to two standard objections to epistemicism. First objection: if vague expressions have sharp cut-off points, how is it that we are not able to say what these cut-off points are? Response: the safety based explanation shows why the cut-off points are unknowable so it is of no surprise that we cannot say what they are. Second objection: if vagueness consists of ignorance, then what makes vagueness a distinctive phenomenon, distinguished from other kinds of ignorance? Response: vagueness involves a special kind of

ignorance, one that is generated because of particular kinds of margin of error principles, namely ones that involve the semantic plasticity of vague expressions.³⁷

I agree with Kearns and Magidor that some explanation must be given if the epistemicist is to respond to the two objections above. If Williamson's explanation is unavailable to the moral realist, are there are other theories that could explain why we can't know the location of the cutoffs and why the ignorance in question is of the special sort that makes for vagueness? Although I cannot conclusively establish that *no* such explanation will be available to the moral realist, I will argue here that the alternative explanations in the literature that have been offered (at least those of which I am aware) are also problematic when extended to moral predicates. I will consider two alternative explanations: one offered by Sorensen and the other by Kearns and Magidor.

Sorensen,³⁸ like Williamson, suggests that vagueness is an epistemic phenomenon. On his view, propositions that ascribe vague predicates to borderline cases (which we'll *indeterminate propositions*³⁹) have definite truth values, and some such propositions are indeed true. However, these truths cannot be known because they lack *truthmakers*; that is, there is no worldly thing that makes these indeterminate propositions true. But this suggestion, as applied to moral vagueness, faces a number of problems.

First, on Sorensen's truthmaker gap view, it is *only* contingent truths that can lack truthmakers due to vagueness. (This is argued for by Mark Jago⁴⁰ and so I will not rehearse the arguments for this here). But moral truths are plausibly necessary.

Second, even if Sorensen's view of indeterminacy could be extended to account for moral truths, indeterminacy due to lack of truthmakers would not be able to explain the unknowability of propositions concerning precise cutoff points. To see this, consider why

Sorensen thinks that, *for contingent* (a posteriori) *truths*, the lack of truthmakers can explain unknowability. Here is what he says:

As Albert Einstein said ‘I am a little piece of nature’. I know contingent truths in virtue of [causal] connections between me and the bits of reality that make those statements true. Absolute skepticism follows when there is no connection between a *contingent* statement and a state of affairs upon which it is contingent.⁴¹ (my emphasis)

Note that Sorensen’s explanation for why we should expect ignorance in the presence of truthmaker-gaps appeals to the fact that the truths in question are contingent. While it may be plausible that I know *contingent* truths in virtue of the causal connections that exist between me and “the bits of reality that make those statements true,” it is not so plausible that this is how, for instance, mathematical or moral truths are known, at least according to realists. Indeed, mathematical and moral realists have been challenged to explain how mathematical and moral knowledge is possible *precisely because* it is presumed that knowledge of the mathematical and moral realm isn’t the result of interaction with “bits.”⁴² Given that truthmakers were never part of the moral realist’s moral epistemology, the absence of such truthmakers won’t (even on Sorensen’s view) explain why propositions concerning precise cutoffs are unknowable.

Third, unlike Williamson’s view, which merely makes knowledge of the sharp cutoff points extremely difficult, or impossible *for agents with our capacities*, the truthmaker gap view supports Sorensen’s claim that these truths are *in principle* unknowable. But, as Dougherty⁴³ and Constantinescu⁴⁴ point out, it’s not clear that there can be moral facts that are in principle unknowable. Constantinescu, relying on a form of internalism about

reasons, claims that if R is a normative reason for X to ϕ , then R must be able to figure in a rational justification for the claim that X ought to ϕ that X could come to know if X's reasoning abilities were maximally improved.⁴⁵ He then argues that if the precise cutoff points of moral predicates are in principle unknowable, moral predicates won't be able to figure into justifications in this way. Dougherty makes a similar point. He cites moral philosophers such as Bernard Williams,⁴⁶ and Judith Jarvis Thomson⁴⁷ who explicitly deny the possibility that there could be unknowable moral facts based on what they take to be the conceptual truth that moral facts must be potentially action guiding. However, it is worth noting that whether there can be in principle unknowable moral facts remains controversial,⁴⁸ so this last argument won't convince those realists who believe in the possibility of moral facts that are in principle unknowable.

The second alternative to Williamson's explanation for unknowability that I will consider is a proposal that Kearns and Magidor⁴⁹ express some sympathy towards, which is based on a view, developed in Breckenridge and Magidor.⁵⁰ On this proposal, vagueness consists in ignorance resulting from the fact that which property a vague predicate refers to is "arbitrary." When we use vague predicates, what essentially takes place is a "reference lottery." Nothing other than, perhaps, luck explains why our term refers to one, rather than another precisification of a vague predicate.

The first thing to note about the view is that, as Kearns and Magidor themselves point out, the view won't help in responding to the second objection to epistemicism: explaining what is unique about the kind of ignorance involved in vagueness. For the paradigmatic cases of arbitrary reference are meant to be cases where no vagueness is involved (such as in cases in which one says: "let 'n' pick out some natural number").

Second, the view will face one of two problems that have already been discussed. If it's in principle impossible to know which property won the reference lottery, then there will be moral facts that are in principle unknowable. On the other hand, if it is *in principle* possible to know which property won the reference lottery, then we face a problem similar to the one that Williamson's view faces: implausible consequences concerning how moral deliberation could be resolved. To see this, note that, on the arbitrary reference view, the referent of "permissible" will differ even amongst very similar linguistic communities. This is because it's highly unlikely that the property that wins the lottery for community A will be the same property that wins the lottery in the neighboring community B. If it's in principle possible to know which property won the reference lottery, then it will turn out that Cheryl could, in principle, resolve her deliberation by finding out about the outcome of the particular lottery that "took place" in her community. However, it's just as implausible to think that one could resolve one's moral deliberation by learning about the outcome of the reference lottery as it is to think that one could resolve one's deliberation by learning about the usage facts in one's community. Indeed, if it turned out that I could resolve my deliberation about whether, say, some abortion is permissible by finding out whether the reference lottery deemed it so, it's hard to see why attribution of moral properties would have any motivating force. If I want to have an abortion, why should the fact that the reference lottery went this way, rather than that way, motivate me to not perform it?

Let's take a step back and consider what the epistemic views of vagueness that I have considered have in common. One interesting feature of such views is that, although proponents of the epistemic conception describe vagueness as an epistemic phenomenon, *the explanation* for the epistemic phenomenon frequently relies on semantic features of

these predicates: features that the moral realist will resist. Williamson requires a shifty semantics and Kearns and Magidor require a semantics in which reference is fixed arbitrarily. However, moral predicates, at least for a realist, simply don't have the necessary semantic features. So it is, at root, for *semantic* reasons that moral realists can't endorse these varieties of epistemicism. The problems with Sorensen's view as extended to moral predicates aren't semantic, but rather metaphysical (it would require moral truths to be contingent) and epistemological (it would require moral truths to be in principle unknowable).

In sum, the challenge for epistemicist accounts of moral vagueness is that the usual ways of responding to certain objections to epistemicism is not available when applied to moral predicates. Thus, if the moral realist is to be an epistemicist, she doesn't *only* have to deal with the general worries about the plausibility of epistemicism. She faces additional problems that an epistemicist about non-moral predicates doesn't face in responding to two central objections to the view.

Let me end by noting that a certain view, which may look like an epistemic account of vagueness, *is* compatible with moral realism. This is the view that, in all of the cases I've been discussing, agents are simply ignorant of whether the act in question is permissible, where this ignorance is not of any special sort. The reason that this is not an epistemic account of vagueness is that it is not an account of vagueness at all. It is an account according to which there is no moral vagueness. This is because mere ignorance doesn't constitute vagueness. We don't know how many stars there are in the universe but this is no manifestation of vagueness. It is for this reason that it has been important for advocates of the epistemic conception of vagueness to describe a kind of ignorance that is *special*.

What we've seen here is that the epistemic views of vagueness currently on offer cannot do this in a way that is consistent with both a plausible semantics for moral predicates and a plausible metaphysics and epistemology for moral facts.

5. Conclusion

There is a sense, it seems, in which it can be arbitrary whether something is red, but not whether something is permissible. This sense of arbitrariness, or lack thereof, is what explains why, for a moral realist, traditional accounts of vagueness aren't well suited for moral predicates.

Traditional supervenient accounts, as well as Williamsonian epistemicist accounts, assume that vague predicates have a *shifty* semantics: minor shifts in usage can affect the referent of the term. This means that questions about, say, redness can sometimes be settled by learning facts about usage. That's all and well and good as far as redness goes. But the realist shouldn't think that questions about *morality* can be settled by learning usage facts. So traditional supervenient accounts and Williamsonian epistemicist accounts are unavailable to the realist. (Other epistemic accounts like Sorensen's and Kearns and Magidor's face related, but distinct worries).

If the semantics for moral predicates isn't shifty, it must be rigid: minor shifts in usage don't shift the referent. I considered two ways that such a view could be developed (reference magnetism and conceptual role semantics). I argued that, on any plausible view of moral vagueness, if the referent of our word is not going to be fixed by factors that the realist considers irrelevant, it is because of the way in which the normative facts constrain the reference of our terms. But these constraints have the result that even omniscient

speakers of a perfect language would encounter moral vagueness, and hence, according to the Barnes criterion, moral vagueness is ontic. It is vagueness that would remain even if we spoke a perfect language and were omniscient. So semantic accounts according to which moral predicates are *rigid* are available to the moral realist, but they commit her to ontic vagueness.

In sum, when we try to apply the standard theories of vagueness to the moral realm, what results is either a view that is unavailable to the realist, or a view that is committed to ontic vagueness. I conclude that if facts about morality are out there in world, whatever vagueness comes along with them will have to be out there as well.

¹ Anthony Ellis, “Deontology, Incommensurability and the Arbitrary,” *Philosophy and Phenomenological Research* 52(4) (1992): 855-875.

² Larry Alexander, “Deontology at the Threshold.” *University of San Diego Law Review* 37 (2000): 893-912.

³ Alastair Norcross, “Comparing Harms: Headaches and Human Lives.” *Philosophy and Public Affairs* 26(2) (1997): 135-167.

⁴ Larry Temkin, “A Continuum Argument for Intransitivity.” *Philosophy and Public Affairs* 25(3) (1996): 175-210.

⁵ Norcross “Comparing Harms,” 156.

⁶ Tom Dougherty, “Vague Value,” *Philosophy and Phenomenological Research* 89(2) (2014): 352-372, 358.

⁷ Elizabeth Barnes, “Fundamental Indeterminacy,” *Analytic Philosophy* 55(4) (2014): 339-362, 339. Barnes actually focuses on the more general phenomenon of indeterminacy, but I will be focusing specifically on vagueness.

⁸ See, for example, Gareth Evans, “Can there be vague objects?” *Analysis* 38 (1978): 208, David Lewis *On the Plurality of Worlds*, (Oxford: Blackwell, 1986), 212 and Mark Sainsbury, “Why the World Cannot be Vague,” *Southern Journal of Philosophy* 33 (1994): 63-82.

⁹ See, for example, Barnes “Fundamental Indeterminacy”, Jessica Wilson, “A Determinable-Based Account of Metaphysical Indeterminacy,” *Inquiry* 56(4) (2013): 535-579, Elizabeth Barnes and Robbie Williams, “A Theory of Metaphysical Indeterminacy in Karen Bennett and Dean W. Zimmerman (eds.) *Oxford Studies in Metaphysics* Vol. 6 (Oxford: Oxford University Press, 2011), Elizabeth Barnes, “Ontic Vagueness: A Guide to the Perplexed,” *Noûs* 44 (4)(2010): 601-627, Robbie Williams, “Ontic Vagueness and Metaphysical Indeterminacy,” *Philosophy Compass* 3/4 (2008): 763-788, and Gideon Rosen and Nicholas J.J. Smith, “Worldly Indeterminacy: A Rough Guide, *Australasian Journal of Philosophy* 82(1) (2004): 185-198.

¹⁰ See note 8.

¹¹ Indeed, Cristian Constantinescu (“Moral Vagueness: A Dilemma for Non-Naturalism” in Russ Shafer-Landau (ed.) *Oxford Studies in Metaethics* Vol. 9 [Oxford: Oxford University Press, 2014], 152-185) has provided a different argument appealing to moral vagueness for such a trilemma.

¹² See Constantinescu, *ibid.* and Dougherty “Vague Value” for a more detailed defense of the existence of moral vagueness.

¹³ Thanks to Ian Proops for this example. Also, I am using “borderline case” here as a neutral term. The existence of borderline cases is meant to be consistent with semantic, epistemic and ontic accounts of vagueness.

¹⁴ Here are two examples of theories on which there will be no moral vagueness: The first is a version of hedonic utilitarianism according to which one is always required to maximize the amount of pleasure in the world, and there is no vagueness or incommensurability concerning which outcomes bring about the most pleasure. The second is a theory that is entirely scalar and so lacks on/off terms like “permissible” or “required.” (Thanks to an anonymous referee for the second suggestion).

¹⁵ The moral realist can allow that *some* of our moral obligations depend on the conventions in our community. For example, if it’s generally impermissible to be rude for no good reason, and rudeness depends on the conventions of a community, the moral realist can say that which actions are permissible depends on the community’s conventions. If you think the diversion case is such a case, substitute one of the other cases. There is no reason to think that, in all cases of moral vagueness, the permissibility facts will be fixed by convention. Thanks to Ulrike Heuer for this point.

¹⁶On the shifty semantic view I have in mind, to the extent that the referent of a term like “permissible” is fixed by our usage, it is not only our individual usage that is relevant, but also the usage of other members of our linguistic community. Diana Raffman (“Vagueness without Paradox,” *The Philosophical Review* 103(1) (1994): 41-74) argues for a view according to which the truth-value of a sentence with a vague predicate depends largely on the *individual’s* judgment at the time the judgment is made. On her view (extended to moral predicates), if I say at time *t*: “an abortion after 180 days is permissible”, the truth-value of my utterance depends largely on whether, at *t*, I *saw* such an abortion, or sufficiently similar ones, as permissible. But the realist will deny that whether or not my moral judgments are correct depends on whether, at the time of judgment, I *saw* the acts in

question as permissible. It is worth noting here that Crispin Wright's view ("On Being in a Quandary," *Mind* 110(437) (2001): 45-97, and "Wang's Paradox" in *The Philosophy of Michael Dummett* [Chicago: Open Court Press, 2007], 414-44) according to which vague predicates put us in a "quandary" is based on a similar form of response-dependence and so is also unavailable to the realist. Sentences with ordinary vague predicates, Wright thinks, have the feature that "their truth or falsity have to be thought of as determined by our very practice, rather than by principles which notionally underlie it" (Wright, "Wang's Paradox," 441). Since the moral realist thinks that the truth or falsity of our moral judgments is *not* determined by our practice – we are not guaranteed that our moral judgments are true, just in virtue of our having made them – Wright's account will also be unavailable to the realist. In fact, Wright explicitly points out that his arguments only support his account when it comes to what he calls "the usual suspects" ("red", "bald", "heap" and so forth). These are predicates that, in his terminology, are "essentially linguistic."

¹⁷ Note that if the shifty account is correct, and the truth value of a sentence like "X is permissible" is highly sensitive to subtle facts about a community's language use, it's not surprising that an agent could be reasonably uncertain about whether a (determinately) permissible abortion is, in fact, permissible.

¹⁸ One might think that there is a way in which learning usage facts could be relevant to deliberation. For perhaps the fact that many members of our community label some abortion as "permissible," could be good *testimonial evidence* for its permissibility. But whether or not this information would be relevant would depend on the correct *epistemic* theory concerning the potential justifying role of moral testimony (an issue that is quite

contentious. See, e.g. Sarah McGrath, "The Puzzle of Pure Moral Deference," *Philosophical Perspectives* 23(1) (2009): 321-344, Paulina Sliwa, "'In Defense of Moral Testimony,'" *Philosophical Studies* 158(2) (2012): 175-195, and Alison Hills, "Moral Testimony," *Philosophy Compass* 8(6) (2013): 552-559. On the shifty semantic account, however, the way in which usage facts would bear on your deliberation has nothing to do with the potential justifying role of moral testimony. What I am claiming is implausible is that learning usage facts could settle deliberation about moral matters in virtue of the fact that the *referent* of moral terms is affected by subtle changes in usage.

¹⁹ Thanks to Matti Eklund for pressing me on this.

²⁰ Ralph Wedgwood, "Conceptual Role Semantics for Moral Terms," *The Philosophical Review* 110(1) (2001): 1-30.

²¹ Dougherty, "Vague Value," 358.

²² David Lewis, "New Work for a Theory of Universals," *Australasian Journal of Philosophy* 61 (1983): 343-377.

²³ Why doesn't the unattractive consequence arise for views according to which the referent of "permissible" shifts with *major* variations in usage? Recall that the problem only arises if learning, say, whether an abortion is in the extension of "permissible" (through usage data), will tell the agent whether the abortion is permissible. But this won't be true if the potential referents for "permissible" that the agent is unsure about span a huge range. For example, imagine that a German speaker, who knows no English, is deliberating about whether an abortion is permissible. If the German speaker found out from the linguistic anthropologist that an abortion isn't in the extension of the word "permissible," this wouldn't resolve her doubts about the permissibility of abortion. After

all, for all she knows, the word “permissible” refers to zebras and finding out that an abortion isn’t a zebra doesn’t tell her anything about whether the abortion is permissible. Thus, when the span of candidate referents is large enough, learning whether an action is in the extension of “permissible” won’t actually tell an agent, who is unsure about the extension of “permissible,” whether an act is permissible.

²⁴ Terry Horgan and Mark Timmons, “Troubles for New Wave Moral Semantics, The Open Question Argument Revived.” *Philosophical Papers* 21 (3) (1992): 153-175.

²⁵ Matti Eklund, *The Language and Metaphysics of Normativity* (working title) (Oxford: Oxford University Press, forthcoming). See also his “Alternative Normative Concepts,” *Analytic Philosophy* 53(2) (2012): 139-157.

²⁶ This is not necessitated by all semantic views of vagueness. For example, according to interest-relative views like those developed by Delia Graff Fara (“Shifting Sands: An Interest Relative Theory of Vagueness,” *Philosophical Topics* 28 (2000): 45-81) and Agustín Rayo (“Vague Representation,” *Mind* 117 (2008): 329-373), the truth of a sentence with a vague predicate depends primarily on the *purposes* for which we’re using the predicate, not the individual instances to which the term is applied by members of the community. If they are right, then perhaps the utterances of any two similarly situated agents, in sufficiently similar linguistic communities, will have the same truth-value in virtue of the shared *interests* of the two agents, and despite the differences in the communities’ applications of the term. Interestingly, however, the mechanisms described by Graff Fara and Rayo, which allow an agent’s purposes to determine the truth-value of a sentence, do not work in the paradigmatic cases for which we use moral predicates. This is because Graff Fara’s and Rayo’s accounts assume that when we use a vague predicate, like “tall,” it is being used as a

tool to rule in or out possibilities that are not themselves specified in terms of tallness. For example, I might use the word “tall” to convey information to you about which of the people in the crowd is my neighbor. However, if I all wanted to convey to you was information about which person was *tall*, my purposes will not help fix the truth-value. The only factors relevant to reference, in such a case, will be the ordinary ones such as usage and naturalness. Graff Fara and Rayo’s accounts work well for many ordinary vague predicates precisely *because* it is rare that what we ultimately care about is whether the predicate itself applies. Usually our interest concerning whether a vague predicate applies can be satisfied by learning information that doesn’t itself contain the vague predicate in question. But moral predicates are different. We are frequently interested in whether something is *permissible*, or whether we *ought* to do it, where this interest is irreducible to an interest in whether some other property of the act obtains. For this reason, the sorts of mechanisms that Graff Fara and Rayo appeal to can’t be used to explain the rigidity of moral predicates.

²⁷ Dougherty, “Vague Value.”

²⁸ Wedgwood, “Conceptual Role Semantics for Moral Terms.”

²⁹ *Ibid.*, 18.

³⁰ Perhaps the conceptual role semanticist can say something along these lines: “Since it turns out that there is no property which determinately makes the rule of practical reason valid, I’ll have to give up on my view that the referent of moral predicates, (even in a perfect language), will be a property that makes the rules of practical reason valid. The best I can hope for is a property such that it’s *indeterminate* whether the property makes the rules of practical reason valid. On my view, a perfect language will contain only

predicates naming precise properties, and the referent of any given moral predicate will be some precise property such that it's indeterminate whether that property makes the rules of practical reason valid." The first thing to note about this suggestion is that it gives up on the conceptual role semanticist's ambition of fixing the referents of moral predicates by appealing to their conceptual role. On the version of the view that I've been discussing in the main text, we don't say that *no* property makes the rules of practical reason valid. Rather, we say that there is some property that makes the rules of practical reason valid but it's indeterminate what it is. More importantly, though, this version of the conceptual role view will no longer be rigid. To see why, note that on the modification we are considering, the conceptual role semanticist hopes to identify the referent of a moral predicate with one of the many properties such that it's indeterminate whether these properties make the rules of practical reason valid. But what is going to determine which, of the many precisifications, the moral predicates of the perfect language users refer to? The most promising development of this proposal would be that use takes over and determines which precisification the community refers to. But then the view is shifty. It has the consequence that using the language slightly differently would result in a different referent.

³¹ Barnes, "Ontic Vagueness," 605.

³² Timothy Williamson, *Vagueness* (Oxford: Routledge Press, 1994).

³³ An example of an epistemic account that is not epistemicist is developed by Crispin Wright ("On Being in a Quandary," "Wang's Paradox"). His account however is unavailable to the moral realist for the reasons discussed in note 16.

³⁴ Timothy Williamson, "Reply to Commentators," *Philosophy and Phenomenological Research* 57(4) (1997): 945-953, 947-948.

³⁵ So named because it relies on the safety requirement for knowledge, roughly: for S to know that p, it must be the case that there are no nearby possible worlds in which S believes p, but p is false.

³⁶ Roy Sorenson, *Blindspots* (Oxford: Oxford University Press, 1988).

³⁷ Stephen Kearns and Ofra Magidor, "Epistemicism about Vagueness and Meta-Linguistic Safety," *Philosophical Perspectives* 22 (2008): 277-304, 277-278.

³⁸ Roy Sorensen, *Vagueness and Contradiction*, (Oxford: Oxford University Press, 2001).

³⁹ I am here following the terminology of Mark Jago ("The Problem with Truthmaker-Gap Epistemicism," *Thought* 1(4) (2013): 320-329). Note that an "indeterminate proposition" doesn't refer to a proposition that is indeterminate in *truth value*. It's just a term used to pick out those propositions that ascribe vague predicates to borderline cases. These propositions, according to Sorensen, have *determinate* truth values.

⁴⁰ *Ibid.*, 325-326.

⁴¹ Sorensen, *Vagueness and Contradiction*, 171.

⁴² See, for example, David Enoch, "The Epistemological Challenge to Metanormative Realism: How Best to Understand It, and How to Cope with it," *Philosophical Studies* 148(3) (2010): 413-438, Hartry Field, *Realism Mathematics and Modality* (Oxford: Blackwell, 1989), 25-30 and Paul Benacerraf, "Mathematical Realism," *Journal of Philosophy* 70(19) (1973): 661-679.

⁴³ Dougherty, "Vague Value."

⁴⁴Constantinescu, "Moral Vagueness."

⁴⁵ Ibid., 179.

⁴⁶ Bernard Williams, "Internal and External Reasons" in Ross Harrison (ed.), *Rational Action* (Cambridge: Cambridge University Press, 1979), 101-113.

⁴⁷ Judith Jarvis Thomson, *Normativity*, (Chicago: Open Court Press, 2008).

⁴⁸ See, e.g. Roy Sorensen, "Unknowable Obligations," *Utilitas* 7(2) (1995): 247-271.

⁴⁹ Kearns and Magidor, "Epistemicism about Vagueness and Meta-Linguistic Safety."

⁵⁰ Wylie Breckenridge and Ofra Magidor, "Arbitrary Reference," *Philosophical Studies* 158 (3) (2012):377-400.