

MORAL VAGUENESS AND EPISTEMICISM

John Hawthorne

EPISTEMICISM is one of the main approaches to the phenomenon of vagueness. But how does it fare in its treatment of *moral* vagueness? This paper has two goals. First, I shall explain why various recent arguments against an epistemicist approach to moral vagueness are unsuccessful. Second, I shall explain how, in my view, reflection on the sorites can inform normative ethics in powerful and interesting ways. In this connection, I shall be putting the epistemicist treatment to work, engaging with a family of somewhat neglected issues concerning continuity that lie at the interface of metaphysics and ethics.

Section 1 introduces epistemicism as well as a competing view—“classical magnets”—that will be helpful for later discussion. Section 2 address a variety of arguments against epistemicist treatments of vagueness in ethics, including Miriam Schoenfield’s appeal to the irrelevance of linguistic anthropology to ethics, Tom Dougherty and Cristin Constantinescu’s concerns about unknowable moral truths, and a proportionality argument leveled by Constantinescu.¹ Section 3 precisifies an interesting but underexplored cluster of continuity issues in the vicinity of the proportionality idea, and examines them through an epistemicist lens.

1. EPISTEMICISM AND CLASSIC MAGNETS

1.1. *Epistemicism*

I will begin with a brief sketch of epistemicism. I shall present the version articulated by Timothy Williamson, who has done the most to popularize the view.² A somewhat barebones version of Williamsonian epistemicism will be adequate to the dialectical purposes of this paper.

1 See Schoenfield, “Moral Vagueness Is Ontic Vagueness”; Dougherty, “Vague Value”; and Constantinescu, “Moral Vagueness.”

2 See Williamson, *Vagueness*. Interestingly, different versions of epistemicism are presented by Sorensen, *Vagueness and Contradiction*; and Kearns and Magidor, “Epistemicism about Vagueness and Meta-Linguistic Safety.” I shall not be discussing their comparative merits and detractions here.

Here are the key features of the epistemicist treatment of borderline cases. First, excluded middle holds in borderline cases. (Epistemicism is one of many theories of vagueness that operates within a classical propositional logic.³) Supposing, for example, that I constitute a borderline case of being happy. Then:

Either I am happy or it is not the case that I am happy.

Second, bivalence holds in borderline cases.⁴ Thus:

"I am happy" is true or "I am happy" is false.

Third, borderline cases beget ignorance among humans, left to their own devices. Thus if I am a borderline case of being happy then humans are not in a position to know whether I am happy.

Some important points of clarification are in order: I say "among humans" since it is no part of this brand of epistemicism that it is impossible in principle for *any* creature to know the answer in a borderline case:

On the epistemic view, vague utterances in borderline cases are true or false and we humans have no idea how to find out which. It is quite consistent with this view that what is a borderline case for us is not a borderline case for creatures with cognitive powers far greater than any we can imagine.⁵

Epistemicism thus allows that borderline cases can be known by superbeings.

I say "left to their own devices" because epistemicism also presumably allows *us* in principle to know the truth value of borderline cases by relying on

3 By this I mean that the epistemicist accepts as true any sentence that is certified as true by the standard truth-table method and accepts as truth preserving any argument that is certified as valid by the standard method of truth tables. Examples of other approaches that endorse classical propositional logic (though this list is far from exhaustive) are the supervaluationism of Fine ("Vagueness, Truth and Logic") and Bacon's version of the view that vagueness resides in propositions and not language (*Vagueness and Thought*). These views also accept the standard inference rules for the existential and universal quantifiers. Not all of these approaches accept bivalence, however. For example, versions of supervaluationism that identify truth with supertruth will accept as supertrue any sentence of the form " x is bald or x is not bald" but will not always accept sentences of the form " x is bald' is true or ' x is bald' is false." The standard way of certifying the latter class as true relies on disquotational principles that will not be accepted by the supervaluationist who goes in for the identification of truth with supertruth.

4 That is, bivalence holds when a borderline case arises on a particular occasion of use. A single, context-dependent sentence may express a truth on one occasion and a falsehood on another, and can be borderline true on one occasion and non-borderline true on another.

5 Williamson, *Vagueness*, 212.

the testimony of superbeings (we might rely on their testimony because we are impressed enough by the performance of the superbeings in areas that we do know about). The relevant ignorance thesis is that the ordinary discriminative and intellectual resources of human beings afford no path to knowledge in borderline cases.

Further, it is tempting to say that if an utterance of “I am happy” is borderline, then, according to epistemicism, the fact it expresses is unknowable. But here we need to be very careful. By way of illustration, let us introduce a name, “Roger,” with the stipulation that “Roger” picks out the number 1 if I am happy and 0 otherwise. The sentence “Roger is 1” is borderline in the way “I am happy” is borderline. But supposing I am happy, it follows that “Roger” refers to 1 and thus arguable that “Roger is 1” expresses the fact that 1 is 1. But that fact is not at all difficult to know! The issue here is that unless we operate with an extremely fine-grained notion of facts, then a fact can simultaneously be expressed by a borderline sentence and also by a precise sentence. In a case like this there will, so to speak, be a blockage to knowing the fact under the guise of the borderline sentence but there may be no similar blockage to knowing the fact under the more precise guise. Reserving “vague” and “precise” as predicates of representations, Williamson introduces the predicates “vague*” and “precise*” as predicates of objects, properties, and relations, so that a vague* object will be one that is picked out by a vague expression and precise* object picked out by a precise expression. He then points out that

The vague description “the greatest prime number much less than 100” and the precise description “the prime number between 72 and 78” might both refer to 73 in a given context. Thus 73 would be both vague* and precise*.⁶

As he makes clear, similar points potentially apply to properties, relations, and facts themselves.⁷

Fourth, not all ignorance is ignorance due to vagueness. Moreover, not all irremediable ignorance is ignorance due to vagueness. For example, certain microphysical or mathematical questions may be deeply elusive but it need not follow that the questions are vague. Williamsonian epistemicism has a distinctive story to tell about the *source* of ignorance in the case of vagueness. The source is semantic plasticity. If a borderline sentence expresses a true proposition, then there is a false proposition that could very easily have been expressed

6 Williamson, *Vagueness*, 258.

7 Of course, this taxonomy allows that there are facts that are vague* but not precise*—in such cases, the only guises under which it is humanly possible to represent the fact are vague guises.

by that sentence on account of tiny differences in use. In this case, the sentence is *true but not determinately true*. And if a borderline sentence expresses a false proposition, there is a true proposition that it could very easily have expressed. In this case the sentence is *false but not determinately false*.⁸

I shall elaborate on some further aspects of epistemicism as needed in the discussion that follows.

1.2. Classical Magnets

Tom Dougherty juxtaposes the Williamsonian explanation of moral vagueness with a second picture. According to this second picture, moral predicates pick out natural kinds, each of which serves as a “reference magnet.”⁹ Small differences in use would not induce a shift in reference because moral predicates referentially gravitate to these properties. Even if other properties “fit our use” a bit better, the predicates refer to the natural kinds because they are much more natural than other properties with similar extensions. He takes the category of being an orangutan as an analogy:

Since we are assuming that there is a natural biological kind, and our usage of the term “orangutan” comes close to picking it out, this natural kind becomes the referent of the word. In this way, a natural kind can act as a “reference magnet” for a term. Let us suppose for now that there is a unique set of things that constitutes the natural kind, orangutan. (Later we will discuss the view that there are overlapping but distinct sets that are equally natural as kinds.) Now, assuming we reject a metaphysical view of vagueness, this natural kind has a precise boundary: everything is in this set or it is not. Therefore, our use of the word “orangutan” would pick out a precise set of things.¹⁰

8 Epistemicists also often introduce determinacy operators like, “It is determinate that,” though (as explained in Fritz, Hawthorne, and Yli-Vakkuri, “Operator Arguments Revisited”), such operators threaten to behave like Kaplanian monsters. Assume for example that “Roger is identical to 1” and “1 is 1” express the same proposition and that we want to say: Determinately 1 is 1 and Not Determinately Roger is 1. Then the determinacy operator will not interact with quantifiers in the usual way. For example, we cannot reason from “Not (Determinately Roger = 1)” to “ $\exists x (x \text{ is Roger and Not (Determinately } x = 1))$.”

For critical discussion of a range of subtleties arising from the semantic plasticity idea, see Hawthorne, “Epistemicism and Semantic Plasticity”; and Yli-Vakkuri, “Epistemicism and Modality.”

9 Dougherty, “Vague Value,” 357. For further discussion of possible applications of reference magnetism in the domain of ethics, see Dunaway and MacPherson, “Reference Magnetism as a Solution to the Moral Twin Earth Problem.”

10 Dougherty, “Vague Value,” 357–58.

The discussion here needs a bit of cleaning up. First, assuming we buy into classical propositional logic (a commitment that, as I have said, is common to many approaches to vagueness, not just epistemicism), the difference between the reference magnetism picture and the semantic plasticity picture is not aptly captured by such claims as “Everything is an orangutan or it is not.” After all, everything is bald or it is not. Given classical propositional logic, whether a term is shifty or stable has nothing to do with an excluded middle. Further, as we have seen, the primary uses of “vague” and “precise” are to representations. We should not be using ideology like “precise set” without saying what that means. Now, as we have seen, the epistemicist has “precise*” and “vague*” at their disposal—but in this sense a set can be both precise* and vague*.

Moreover, we should think of referents of moral predicates as properties and not extensions (i.e., the sets of things that predicates are true of). On the natural way of developing the magnets picture, “acted permissibly” would have expressed the same property at nearby worlds. But it need by no means have the same extension. Suppose for example that someone behaves permissibly at a dinner party at this world but not at a nearby world. That will suffice to induce a shift in extension of “acted permissibly” across worlds. But this is not the kind of shiftiness that interests the epistemicist.¹¹

Finally, one should not simply assume that insofar as a moral predicate picks out a natural kind, then, even assuming there is no other natural kind with a similar extension, it *follows* that small differences in use of that predicate will be semantically sticky across the modal neighborhood. Consider, for example, a modal path from a world where “is morally good” picks out the property of being morally good to a world where “is morally good” expresses the property of being made of gorgonzola cheese, and where adjacent worlds on the path are almost the same in the distribution of microphysical properties. No matter how powerful a reference magnet moral goodness is, there will be a pair of adjacent worlds along the path, one at which “morally good” picks out moral goodness, and one at which it does not. (Arguably the extension will shift radically between that pair—use will become gradually more and more anomalous vis-à-vis the behavior of moral goodness so that eventually there is such a great mismatch that reference will shift quite dramatically to something much less anomalous. Of course I am not saying that it will jump straight from the property of moral goodness to the property of being made of gorgonzola.) The kind of reference magnetism afforded by natural kinds is one that may secure a good deal of stability but will not in general preclude the plasticity phenomenon and thus will not in full generality preclude the

11 Relatedly, I find the talk of sets as “constituting” natural kinds rather odd.

kinds of shiftiness that Williamson envisages. If the kind picked out fits use just well enough to count as the referent but still badly enough to not fit use well enough at close worlds to count as the referent, then the term will be shiftiness in its referent.

That said there is a picture implicit in Dougherty's sketch (albeit a very speculative one). One can think that a certain moral predicate, say "being morally permissible," picks out a natural, "jointy" property and be anti-skeptical enough to think that we are not in one of the edge cases described in the last paragraph, so that our local modal neighborhood is one in which "is morally permissible" refers to exactly the same property no matter which world from the local neighborhood is actualized. Let us call this view the "classical magnets" view of a moral predicate (where "classical" serves as a reminder that the background propositional logic is classical): we might combine classical propositional logic and bivalence with the idea that at least some moral predicates correspond to highly natural properties, which they stably refer to across nearby worlds.¹² Such predicates are not semantically plastic.

What bears emphasis is that, from the point of view of Williamsonian epistemicism, this is *not* (*pace* Dougherty) a competing explanation of moral vagueness. For given Williamsonian epistemicism, this is a view on which the relevant moral predicates are not vague at all. For recall that not all irremediable ignorance is ignorance due to vagueness. Consider phenomenal consciousness. There may be certain creatures that humans are unable, left to their own devices, to recognize as conscious even though they are. But that does not mean that the question of whether such creatures are conscious is vague. Many of us will be inclined to think that there is a fundamental phenomenon here—being phenomenally conscious—and our ignorance is simply a matter of not having the epistemic tools to probe whether the phenomenon is exhibited by the creature in question. Here a classical magnets view is tempting—it is not unnatural to think that we philosophers have done enough to see to it that "being phenomenally conscious" locks on to a kind, and that tiny differences in use *would not* have induced a semantic shift. I do not care at the moment whether this picture is correct. What I do wish to emphasize is that *if* it is correct, then insofar as ignorance due to vagueness is rooted in semantic plasticity, we should not classify our irremediable ignorance about whether a certain creature is conscious as ignorance due to vagueness.

12 A variant that I shall mention but not discuss holds that the relevant moral properties are magnetic in the sense of being easy to refer to but denies that they are highly natural and thus takes such properties to be a counterexample to Lewis's idea that naturalness correlates with easiness to refer to.

The case of moral vagueness is no different. Let us take what seems like a paradigm case of moral vagueness, one that is presented by Schoenfield:

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

This certainly *seems* like a case of vagueness. But if we fully espouse the classical magnets view for “is morally permissible,” it is by no means clear that we should classify the case in this way. Now of course the sentence, “It is permissible for Darryl to divert his attention for twelve seconds,” may be vague for reasons having nothing to do with moral predicates: after all, there may be actions that are borderline cases of diverting attention. But assuming classical magnets is the right view for “permissible” but not for “bald” and “heap” and the other paradigms of vague predicates, it does not seem, on reflection, that the Williamsonian should see this as a case where “It is permissible that” is vague. Rather, assuming the metaphysically ambitious picture of permissibility encoded by a classical magnets view, this seems relevantly similar to the consciousness case. Just as we cannot discern the length of a rod to the nearest nanometer given our limited discriminatory perceptual capacities, so we cannot see the distribution of the special permissibility property given our limited discriminatory intellectual and discriminatory capacities. Moreover, if there is a single highly magnetic property of propositions picked out by “It is permissible that,” then it does not seem likely that small shifts in use will induce semantic shifts and so the case will not have the semantic plasticity that Williamsonian epistemicism requires for ignorance due to vagueness. What may have seemed like ignorance due to vagueness thus may turn out, given the classical magnets view, to be plain old ignorance.

Two final points of clarification. First, while the Williamsonian version of epistemicism contends that ignorance due to vagueness is rooted in semantic plasticity, I leave open the possibility of a version of epistemicism that locates the distinctive source of ignorance due to vagueness elsewhere. And for all I have said, such views may offer a different take on how classical magnets relate to vagueness. Here is not the place to explore in detail what other epistemicist stories might be available about the distinctive source of ignorance due to vagueness.¹⁴

13 Schoenfield, “Moral Vagueness Is Ontic Vagueness,” 262.

14 I note though that another well-known version of epistemicism on the market does not fit well either with the view that moral predicates are both vague and correspond to classical magnets. Here I have in mind Roy Sorensen’s version of epistemicism, according to

Second, I note that, later in his paper, Dougherty presents a view that relaxes the assumption that moral predicates are magnetized to natural kinds:

Instead there may be multiple natural extensions that are on a par with respect to naturalness.... One might think that the term “orangutan” importantly picks out something special about the metaphysical structure of the world, independently of how we represent it. But one might deny that there is a single precise set of creatures that forms a natural kind. Instead, there may be multiple equally natural sets of creatures that slightly vary in their membership.¹⁵

Adapted to the language of properties, and transposed to the moral case, the idea would be that there is a cloud of very natural properties in the vicinity of permissibility, none of which “is permissible” stably refers to.¹⁶ This way of trying to do justice to the idea that permissibility is metaphysically special, one that denies the stability thesis of the classical magnets view, is of course much more friendly to epistemicism.

2. SOME ANTI-EPISTEMICIST ARGUMENTS

2.1. *Moral Vagueness and Linguistic Anthropology*

Miriam Schoenfield argues for the following thesis:

Central Thesis: If a robust form of moral realism is true, and there is moral vagueness, then it is ontic vagueness.¹⁷

which the distinctive feature of vague truths is that they are truths without a truthmaker (*Vagueness and Contradiction*). Classical magnets seem paradigmatically well suited to provide truthmakers. A referee wondered whether a highly natural cutoff on a scale of moral worth might reasonably be classified as vague. Here again the crucial issue is what the relevant epistemicist story is about the source of ignorance due to vagueness as opposed to ignorance not due to vagueness. If the story about what constitutes ignorance due to vagueness is Williamsonian, then insofar as that natural cutoff induced a failure of semantic plasticity, the ignorance would not count as ignorance due to vagueness. And if the story is Sorensen's then insofar as that natural cutoff provided a truthmaker, then once again the ignorance would not count as ignorance due to vagueness.

¹⁵ Dougherty, “Vague Value,” 364–65.

¹⁶ Here, “is permissible” expresses a property of properties (i.e., action types), “It is permissible that” expresses a property of propositions (a propositional operator), and “acted permissibly” expresses a property of people. I shall not fuss about which member of this family is more fundamental.

¹⁷ Schoenfield, “Moral Vagueness Is Ontic Vagueness,” 259.

By “moral realism” she tell us that she means the view that “moral truths are necessary” and that they are “part of the deep underlying metaphysical structure of the world.”¹⁸ By “moral vagueness” she has in mind the thesis that moral predicates are vague. Focusing on moral permissibility, she presents a series of examples that constitute a *prima facie* case for moral vagueness. We have already looked at one of them—Diversions.

Let us suppose this situation described in Diversions obtains and the following claim is true:

It is permissible for Darryl to divert his attention from his daughter for one second.

Is this a moral truth? One would have thought so except that Schoenfield has told us, on behalf of the moral realist, that she intends to restrict “moral truth” to “necessary truths.” The proposition expressed by the above claim is not necessary, since it would be false in certain circumstances where Darryl learned that his daughter will be tortured if he diverts his attention for a second or more. Now of course there are necessary truths in the neighborhood. Such necessary truths include something of the form:

If *c* then it is permissible for Darryl to divert his attention from his daughter for one second

where “*c*” is a placeholder for an enormously complicated description of the underlying physical facts of the scenario in question (together with phenomenal facts if one thinks of those as a metaphysical add-on).

Necessary truths in the vicinity also include:

Actually, it is permissible for Darryl to divert attention from his daughter for one second.

(Here I am using “actually” in the way that is standard in philosophy:¹⁹ it has a rigidifying effect, so that if “*S*” is in fact true then “Actually *S*” is necessarily true. We can think of claims that a proposition is actually the case as a claim that *at the actual world* that proposition is the case.) But I do not see much point in restricting the category of moral truths to necessary ones.

In a borderline case, the epistemicist subscribes to what Schoenfield calls the “Shifty View”—namely, that “the truth-value of an utterance: ‘*X* is permissible’ is highly sensitive to the way the word ‘permissible’ is used in a linguistic

18 Schoenfield, “Moral Vagueness Is Ontic Vagueness,” 259–60.

19 A standard source here is Davies and Humberstone, “Two Notions of Necessity.”

community.”²⁰ I think what Schoenfield means here is the thesis that the truth value of *certain utterances* of the form “*X* is permissible” is so sensitive. If “It is permissible for Darryl to divert his attention for nine seconds” was highly sensitive but “It is permissible for Darryl to divert his attention for one second” was stable in truth value across slight shifts in use, then the shifty view, as Schoenfield intends it, would be vindicated.

Here is Schoenfield’s master argument against the Shifty View:²¹

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.

Note that the claim that Cheryl can learn what is permissible by crunching through the data does not 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

20 Schoenfield, “Moral Vagueness Is Ontic Vagueness,” 264.

21 One might have expected Schoenfield to say that if “permissible” is shifty, then permissibility is not metaphysically special in the way that the robust moral realist supposes, and hence that the shifty view is unavailable to the robust moral realist. However, she is aware of the “more relaxed” view just described in Dougherty, one that seems to combine plasticity with the thesis that permissibility is highly natural, which perhaps makes permissibility special enough for the tastes of the robust moral realist, even if it is a denizen of a highly natural cloud of candidate referents. (See “Moral Vagueness Is Ontic Vagueness,” 270.)

Note that she thinks the argument that follows cannot be generalized into an argument against shifty treatments of any predicates since “nobody agonizes about borderline cases of baldness.” (See “Moral Vagueness Is Ontic Vagueness,” 266.)

some abortion whose permissibility she was uncertain about, is, in fact (determinately!) permissible by collecting linguistic data. However, it does not 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.²²

The concern, in short, is that the Shifty View allows linguistic anthropologists to resolve moral conundrums by collecting data about linguistic usage.

The Shifty View is (at least approximately) the thesis of semantic plasticity for moral predicates. And as we have seen, the epistemicist subscribes to it. Recognizing this, Schoenfield's complaint against Williamsonian epistemicism is that it is vulnerable to the complaint articulated in the text above:

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.²³

This complaint involves some important misunderstandings that Williamson goes to some lengths to ward off. If our ignorance of the truth value of a borderline use of "I am thin" is the sort of thing that could be resolved by an anthropological investigation into usage, then that ignorance would run no deeper than our ignorance about the relevant facts of usage. But it is crucial to Williamson's epistemicist vision that the ignorance does run deeper. While the epistemicist is very much open to the metaphysical thesis that the intension of, say, "thin" supervenes on various underlying physical facts, including facts about usage, it is crucial to his vision that details of this dependence are epistemologically elusive. In this connection, Williamson points out the metaphysical dependence encoded by supervenience claims does not mean that we could somehow be in a position to know some supervenient fact once we knew some facts on which the former supervene. Against the thought that "since the supervenience generalizations are metaphysically necessary, they can be known a priori," he writes that

as Kripke has emphasized, it is fallacious. Indeed, metaphysical necessities cannot be assumed to be knowable at all.²⁴

Later he writes that

22 Schoenfield, "Moral Vagueness Is Ontic Vagueness," 265–66.

23 Schoenfield, "Moral Vagueness Is Ontic Vagueness," 267.

24 Williamson, *Vagueness*, 203.

one should not be surprised that the known supervenience of *A*-facts on *B*-facts does not provide a route from knowledge of *B*-facts to knowledge of *A*-facts.²⁵

Part of the point is that even if, say, thinness supervened on physical dimensions, knowing physical dimensions would not always put one in a position to know whether someone is thin. But these remarks are also very relevant to the idea, implicit in Schoenfield's discussion, that since the meaning of vague words depends on the sort of facts recoverable by linguistic anthropology, the question about whether, say, it is permissible for Darryl to divert attention for thirty-one seconds can be resolved by linguistic anthropology.

Williamson also points out that the slogan that meaning supervenes on use neglects "the environment as a constitutive factor in meaning."²⁶ But he does not wish to rest everything on that point. His idea is that even granting that some refinement of that slogan is correct (one that dispenses with a crude notion of "use"), it would not vindicate the thought that facts about meaning are somehow accessible to humans. Speaking of the imagined refined gloss on "use" he says:

Although meaning supervenes on use there is no algorithm for calculating the former from the latter. Truth conditions cannot be reduced to statistics for assent and dissent.²⁷

Consider material conditionals about Darryl, with a borderline sentence about permissibly diverting attention inserted as the consequent and with some complicated antecedent giving precise data about usage facts and precise data about Darryl's physical situation. On the epistemicist picture, even if some such an antecedent is determinately true and even if the conditional is necessary, it by no means follows that the conditional is determinately true. Consistent with the conditional expressing a necessary truth there might be a proposition easily meant by the consequent such that the conditional is false when the consequent is interpreted this way. (Thus such conditionals pose no obvious threat to the claim that if the antecedent of a conditional is non-borderline true and a conditional is non-borderline true then its consequent is non-borderline true.)

It is thus crucial to the epistemicist vision that linguistic anthropologists *cannot* resolve borderline issues in the way that Schoenfield envisages. Her objection depends on a misunderstanding of the view.²⁸

25 Williamson, *Vagueness*, 204.

26 Williamson, *Vagueness*, 206.

27 Williamson, *Vagueness*, 206.

28 A referee reasonably wondered whether something in the vicinity of Schoenfield's concern could be revived for the kind of epistemicist—like Williamson—who allows that a

There is one further gap in Schoenfield's argument. Recall from our earlier discussion that it is best to think of candidate interpretations of predicates as properties but not sets. (Similarly, one should think of candidate interpretations of sentences as propositions, not truth values.) Pretend now that a linguistic anthropologist could somehow discern which property was picked out by "is permissible" on an occasion of use. Even then it would not straightforwardly follow that Cheryl's conflict would be epistemically resolved. Knowing which property is expressed by a predicate is one thing; knowing whether a certain object, act, or event instantiates it is, on the face of it, quite another. What Schoenfield's argument thus requires is not merely that the anthropologist can resolve which property "is permissible" picks out but, moreover, that they can resolve this in a way that somehow automatically answers questions about the extension of that property.

Before moving on, I would like to draw attention to one further issue. In the quoted passage, Schoenfield imagines Cheryl and her partner "spending a great deal of time deliberating" and indeed "agonizing" about whether abortion is permissible in the case she describes. Assume it is a case of vagueness. What kinds of attitudes would the epistemicist recommend? Insofar as the case is known by all parties to be borderline it would in some ways be a bit odd to spend a great deal of time deliberating. Normally when one deliberates

superbeing could in principle know the facts of reference. The referee wrote: "The thought I take it is that it's implausible that these sorts of facts (processed by a human anthropologist or a super-being anthropologist) could serve as basis for determining the truth value of 'this abortion is permissible.'" Let us think this through. Williamson's superbeing will be aware of the constitutive dependence of meaning on the environment. Moreover, epistemicists will be very sympathetic to the idea that, even at the use end of things, tiny, inscrutable micro-differences in use that are not fully captured by ordinary anthropological data may make a difference. Given this, our superbeing is going to need a lot more to work with—even on the "use" side of things—than a folder of use facts of the sort that might be cataloged by human anthropologists. So our superbeing will need to know everything relevant about the environment—perhaps up to some astonishing detail, and will likely have to know incredibly fine-grained facts about use. Perhaps the best way to imagine our superbeing, then, is as one who knows the true function from microphysical distributions to facts about reference and is also capable of surveying the microphysical lay of the land in full detail. Our superbeing says "I've surveyed the microphysical landscape, applied the algorithm and determined that the predicate 'is permissible' expresses a property that applies to the referent of 'this abortion.'" But in this setting the thought that the superbeing's reflections could not serve as a basis for determining whether the abortion is permissible has no bite. Cheryl will have no problem knowing that the abortion is permissible iff the property expressed by "is permissible" is instantiated by the event picked out by "this abortion." So if Cheryl is convinced that the superbeing knows exactly how the facts of reference supervene on microphysics and that the superbeing knows the microphysical lay of the land, then she will of course regard the superbeing as having settled the question she is worrying about.

at length about a question it is because one hopes to know the answer. If one knows a case is borderline, one knows one will not find out the answer. So why the prolonged deliberation? That is not to say that the epistemicist recommends a do-not-care attitude. Cheryl can be *concerned* that the action is impermissible. She can *fear* that the action is impermissible. She can have *nonzero credence* that the action is impermissible. But the point of extended deliberation is less clear when one knows that one is not going to discover an answer at the end of it.

2.2. Unknowability

Constantinescu and Dougherty both raise concerns to the effect that the epistemicist approach to moral vagueness delivers unknowable ethical truths, concerns that are cited with guarded approval by Schoenfield.²⁹ We have already seen that epistemicism's purported implication of unknowable moral facts is not incontestable, since a truth that is unknowable under the guise of one vague predicate may, for all epistemicism says, be knowable under another. But I shall look past that point in the following discussion.

Dougherty writes:

How many cents are you required to spend on a taxi in order not to be late for an appointment for which you have promised to be punctual? A friend of an epistemic position may say there is a precise minimum here. But it stretches the imagination to think that we could know what this minimum is.³⁰

He goes on to worry that it is arguably a conceptual truth about ethical facts that they must be action-guiding,³¹ that the postulate of unknowable ethical truths threatens to clash with that conceptual truth.

Cristian Constantinescu discusses epistemicist treatments of moral vagueness and argues that it is incompatible with a nonnaturalist position that takes moral facts to be "intrinsically reason-giving."³² His focal concern is that, while the phenomenon of unknowable facts may be unproblematic, there is something problematic about unknowable ethical facts. The picture here, as Constantinescu develops it, is that every ethical fact constitutes a normative reason for or against certain kinds of actions.³³ He then argues that there is something incoherent

29 Schoenfield, "Moral Vagueness Is Ontic Vagueness," 279.

30 Dougherty, "Vague Value," 361.

31 Dougherty, "Vague Value," 361.

32 Constantinescu, "Moral Vagueness," 155.

33 I think it is useful here to distinguish facts that are naturally expressed using explicitly ethical language and facts that are not so expressed but that have (in some cases contingently) normative significance. That *S* is on fire may well be a reason to help *S* but the proposition

about the idea that a normative reason could be unknowable, since normative reasons have to be in principle available to us as things to justify actions:

But what doesn't seem possible is to divorce *n*-reasons [his shorthand for "normative reasons"] even from a maximally improved capacity for practical rationality. Thus, we can of course accept that there may be moral reasons for us to desist from some of our current practices, but that those reasons are inaccessible to us, due to certain biases or errors in our judgement of which we are unaware. But to recognize them as reasons means to accept that they would serve as justifications for us if our reasoning abilities were improved. What seems incoherent is the thought of an *n*-reason entirely divorced even from the sound exercise of a maximally improved capacity for practical reasoning. To claim that there are reasons which couldn't be anyone's reasons seems almost vacuous. I shall express this upshot in the form of the following epistemic constraint on normative reasons:

Epistemic constraint on reasons: If *R* is an *n*-reason for *X* to ϕ , then *R* can feature in a rational justification of the claim that *X* ought to ϕ , a justification which *X* knows or could come to know if *X*'s reasoning abilities were maximally improved.³⁴

An important thing to notice is that it is not the full apparatus of epistemicism but standard inference rules for the quantifiers and logical connectives, in combination with some fairly banal truths, that delivers the conclusion of unknowable moral truths. Let us use Darryl as our focal example. Take the banal truths

- (a) Darryl is permitted to divert his attention for one second; and
- (b) It is not the case that Darryl is permitted to divert his attention for three hundred seconds.

that *S* is on fire is not something expressed using normative language. My remarks can be adapted to both an expansive and restrictive conception of ethical facts, where the expansive conception includes facts of ethical significance that are not expressed using ethical language.

34 Constantinescu, "Moral Vagueness," 178–79. He then cites with approval Parfit, who wrote: "when it is true that we have decisive reasons to act in some way, this fact makes it true that if we were fully informed and both procedurally and substantively rational, we would choose to act in this way" (*On What Matters*, 1:63). If "full information" includes being fully informed about what is permissible and what is not, then this remark does not in fact lend any support at all to the claim that all ethical truths are knowable. I have not found a consensus among my informants as to what Parfit intended here. Even if "full information" included only the entire body of empirical facts, this would hardly give much support to the claim that all normative reasons are knowable since the relevant empirical facts that figure in the idealization may include unknowable ones.

We can then use standard inference rules to generate a reductio of “For all numbers of seconds n (Darryl is permitted to divert his attention for n seconds $\supset$ Darryl is permitted to divert his attention for $n + 1$ seconds)” in order to derive:

- (c) It is not the case that for all numbers of seconds n (Darryl is permitted to divert his attention for n seconds $\supset$ Darryl is permitted to divert his attention for $n + 1$ seconds).

We can then exploit the duality of universal and existential quantification to reach:

- (d) There is a number of seconds n such that Darryl is permitted to divert his attention for n seconds, but it is not the case that Darryl is permitted to divert his attention for $n + 1$ seconds.

Here is a further truth that also seems fairly banal:

- (e) There is no number of seconds n such that we (Darryl or any other human) are in a position to know that (Darryl is permitted to divert his attention for n seconds but that it is not the case that Darryl is permitted to divert for $n + 1$ seconds).

(If one happens to have some optimism here, then run the whole argument using milliseconds or nanoseconds—here the claim corresponding to (e) is even more secure.)

But (d) and (e) entails:

- (f) There is a number of seconds such that (Darryl is permitted to divert his attention for n seconds and it is not the case that Darryl is permitted to divert his attention for $n + 1$ seconds) and we are not in a position to know that (Darryl is permitted to divert for n seconds and that is not the case that Darryl is permitted to divert for $n + 1$ seconds).

But this is the very kind of claim that enemies of unknowable ethical truths balk at. (Notice that the argument does not deploy words such as “determinately,” “precise,” and “sharp,” so their semantic contribution to sentences in which they occur is neither here nor there.) Every view of vagueness that accepts the relevant banal claims (a), (b), and (e), and the validity of the relevant inferences is committed to conclusions like (f). Such views (as noted earlier—see note 3) include not only epistemicism but supervaluationism, among others.³⁵

35 Williamson makes a point of emphasizing that supervaluationism is committed to claims that encode the idea that tiny differences sometimes make a difference: “Many people have found the major premise [of the sorites] implausible just because it seemed to them that

Dougherty does not seem to be suggesting that we give up claims like the fairly banal (a), (b), and (e). He thus seems to be suggesting in effect that there is a compelling ethical argument against standard inference rules for the quantifiers. Turning to Constantinescu, the following claims seem extremely plausible (at least insofar as we are comfortable with the ideology of normative reasons):

- (g) For any n , if it is permissible for Darryl to divert attention for n seconds, then the fact that it is permissible for Darryl to divert attention for n seconds is a normative reason.
- (h) For any n , if it is impermissible for Darryl to divert attention for n seconds, then the fact that it is impermissible for Darryl for n seconds is a normative reason.

Consider now

Normative-Reasons Conjunction:

- (i) $(p \text{ and } p \text{ is a normative reason}) \text{ and } (q \text{ and } q \text{ is a normative reason}) \supset ((p \text{ and } q) \text{ and } ((p \text{ and } q) \text{ is a normative reason}))$.³⁶

This principle is obligatory for those who hold that all ethical truths are normative reasons since it is obvious that the conjunction of any pair of ethical truths is an ethical truth. But it is plausible in its own right. But putting (g) together with (h), (i), and (j), we can conclude:

- (j) There is a number of seconds n such that (i) Darryl is permitted to divert his attention for n seconds but it is not the case that Darryl is permitted to divert his attention for $n + 1$ seconds, (ii) that conjunctive fact about n is a normative reason, but (iii) we are not in a position to know that conjunctive fact.

Thus, in effect, Constantinescu seems to be suggesting that the epistemic constraint on reasons provides good grounds for rejecting standard inference rules for the quantifiers and connectives.

Giving up the relevant inference rules strikes me as something of an overreaction to examples like Darryl. At any rate I invite readers to consider whether they wish to go in that direction and, if so, what inference rules they recommend putting in their place. Certainly these authors do not suggest any alternative logical framework in which to evaluate the claim that there are unknowable

there could not be a number n such that $n + 1$ grains make a heap and n do not. Supervaluationism makes the very claim that they find incredible" (*Vagueness*, 153).

³⁶ In saying this I do not mean to suggest that *strength* of reasons can be computed in a flat-footedly additive way.

moral truths.³⁷ This suggests to me that they have not really confronted the choice between giving up the relevant inference rules and accepting unknowable moral truths.

But what of Constantinescu's purported connection between normative reasons and moral facts and Dougherty's purported connection between motivating reason and moral facts? Dougherty's view can be put in its best light if we accept the following principle connecting reasons to action:³⁸

Knowledge Principle for Personal Reasons (KRP): If p is S 's reason for ϕ -ing, then S knows p .

At least on the most natural way to resolve the notoriously flexible possessive construction, this principle has a good deal going for it.³⁹ At any rate, I shall assume it here, which helps rather than hinders Dougherty's concerns. A similar idea is in play in Constantinescu, since he relies on the idea that "To claim that there are (normative) reasons which could not be anyone's reasons seems almost vacuous."⁴⁰ If we supplement this thought with the thought that for p

37 Note that a flat-footed, tripartite division of propositions into true, false, and neither does not on its face help much, as it is hard to imagine that belief in truths on the boundary will be safe enough to be known.

Of course, there are some in the Dummettian tradition that self-consciously try to preserve the knowability of moral truths by dispensing with classical logic. (See for example Wright, "Ethical Truths.") I do not have the space here to get into larger questions about the level of abductive support that is enjoyed by classical propositional logic and the standard inference rules for the quantifiers that are recommended by classical model theory. However, it should be quite obvious to readers that I hold these in high regard. Intuitionists reject the inference from the negative universal, "It is not the case that every number n is such that if n grains make a heap, n plus one makes a heap" to an existential conclusion. (Fine opts for an even weaker logic than the intuitionist one, one that precludes inferring Q from $(P \text{ and Not } (P \text{ and Not } Q))$; "The Possibility of Vagueness.") What is striking, though, is that many of the writers in the ethics literature who raise knowability worries—Dougherty and Constantinescu being paradigms—are not mindful of the threat to the standard inference rules and certainly do not offer a competing logic as a working alternative.

For another defense of unknowable moral truths, see McGrath, "Moral Realism without Convergence."

38 For defenses of the idea encoded by KRP, see Dietz, "Reasons and Factive Emotions"; Unger, *Ignorance*; Hawthorne and Magidor, "Reflections on the Ideology of Reasons"; and Williamson, "Acting on Knowledge." The expression "personal reason" is borrowed from Grice, *Aspects of Reason*.

39 For a discussion of the different ways to read possessive constructions in this kind of context, see Finlay, *Confusion of Tongues*, chap. 5.

40 Constantinescu, "Moral Vagueness," 178–79.

to be someone's reason it has to be known, we get his desired conclusion, but without that supplement it is hard to see how to reach that conclusion.

Even granting KRP, it will seem excessive to many to give up standard inference rules for quantifiers and/or connectives on the basis of the lines of thoughts these authors advance. Nevertheless, it may be helpful to offer further therapy to those who are still tempted by them.

It is important to notice that there is something very misleading about the thought that if a proposition is unavailable as a personal reason then it cannot guide action. I do not want to fuss too much here about the term of art "guiding action." But—to return to a theme from the end of section 1.3—it bears emphasis that one can stand in all sorts of interesting relations to an unknowable fact, relations that can certainly have a bearing on one's planning and one's action. One can know in many cases that there is a chance that *p*. One can in some cases know that there is a significant chance that *p*. One can have significant moral concern that *p*. One can take precautions against *p*. One can fear that *p*. One can have a decently high rational credence that *p*. And so on. So the thought that unknowable propositions are dead to one as far as actions and planning are concerned does not seem to be a very good thought. Of course, assuming the personal reasons principle, there is one way that an unknowable proposition *cannot* be a guide to action—namely, by serving as a personal reason for action. But it seems extremely unpromising to elevate this relation over all others in one's account of ethical facts and not much more promising to elevate this relation over all others in one's account of what can stand as a normative reason.⁴¹

2.3. *A Proportionality Argument against Epistemicism*

Constantinescu has another argument that I think is unpersuasive but that (as we shall see in the next section) points toward some interesting issues.

The motivating concern is that, according to the epistemicist, incredibly small differences can make the difference between doing something permissible and not doing something permissible. Schoenfield's case of Darryl is adequate for our purposes here (Constantinescu discusses a very similar case): there are two periods *p*₁ and *p*₂ only a nanosecond apart such that it is permissible for Darryl to divert his attention for *p*₁ but not for *p*₂. In a case like this, the line of thought runs, an idealized agent who was apprised of what was

41 While I do not wish to put too much stock on the point, we should also recall that it is not really part of epistemicism that it is impossible *simpliciter* for humans to know borderline claims. They could, for example, learn them by testimony from superbeings. Suppose a superbeing knew a certain borderline diversion was impermissible for Darryl. The superbeing might well offer some advice: "Don't do that! It is impermissible!" Would this not put Darryl in touch with a normative reason?

going on would have praised Darryl in the one case but would have “blamed and chastised” him in the other.⁴²

Constantinescu worries that “something is amiss. . . . The slightest difference (one nanosecond, one nanogram, one nano-anything) is all it takes for an option to change moral valence. This appears to violate considerations based on justice.”⁴³

The line of thought turns on something like the following inconsistent triad:

1. Other things being equal, very different reactions/treatments of agents are warranted as between any case *C*₁ where Darryl diverts attention from his child in an impermissible way and a case *C*₂ where Darryl diverts attention in a permissible way. (I say “other things being equal” simply to control for pairs of cases *C*₁ and *C*₂ where Darryl acts permissibly in *C*₁ as far as diverting attention goes but commits some other sin that is not committed in *C*₂.)
2. But if epistemicism is right there is a pair of cases incredibly close together microphysically—indeed where the difference in attention is only one nanosecond apart—where Darryl acts impermissibly in one case and permissibly in another and where other things are equal.
3. If cases are almost microphysically the same they warrant almost the same reaction/treatment.⁴⁴

Constantinescu evidently thinks 2 is the culprit, suggesting that “we should reject the epistemicist’s idea of sharp properties on moral grounds.”⁴⁵

As noted earlier, we should be cautious of applying predicates like “sharp” and “precise” to properties. Such predicates, for the epistemicist, apply to representations. We can contrive “vague*” and “precise/sharp*” along Williamsonian lines, but then we should remember that properties can be both vague* and sharp*. What Constantinescu is getting at with his own use of “sharp” is, I think, merely the epistemicist’s commitment to classical propositional logic and standard rules for the quantifiers. And indeed these logical commitments

42 Constantinescu, “Moral Vagueness,” 180.

43 Constantinescu, “Moral Vagueness,” 181.

44 At one point Constantinescu goes so far as to say that if people *X* and *Y* are almost indistinguishable then it would be unjust for them “to receive different treatment” (“Moral Vagueness,” 181). The idea that pairs of people that are almost indistinguishable should not be treated differently *at all* is indefensible. Pairs of cases that are almost the same can be chained together so that there are cases wildly different at each end. But repeated application of the principle that almost indistinguishable cases cannot be treated differently at all would have us conclude that the cases at each end of the series call for the same treatment.

45 Constantinescu, “Moral Vagueness,” 180.

alone, with or without the extra commitments of epistemicism, get us the conclusion that he finds abhorrent. Grant that it is permissible for Darryl to divert his attention for one second but impermissible to divert his attention for five minutes. Armed with the relevant logical tools we can simply prove that there are two periods of time, P_1 and P_2 , one nanosecond apart, such that it is permissible for Darryl to divert his attention for P_1 but that it is not the case that it is permissible for Darryl to divert his attention for P_2 . We do not need to make any use of such ideology as “sharp properties.” As before, what Constantinescu is in effect telling us is that we have decisive moral grounds to reject the relevant inferences rules! Again, this will seem excessive to many. And, as before, he offers no alternative logic and so, for example, we are left in the dark about how we are supposed to reason with any of the principles in the paper that he is friendly to.

As an argument against classical propositional logic and/or the standard inference rules for the quantifiers, I am completely unmoved. We should all make our peace with the sorites and realize that, for pretty much any predicate, tiny differences sometimes make a difference between the predicate being true of a thing and being false of a thing. But the inconsistent triad is not without interest. Even if we dismiss the line of thought as grounds for logical deviance we are left with the interesting decision as to whether to give up 1 or 3 in the inconsistent triad.

In support of 3, Constantinescu offers a sweeping claim about supervenience:

Proportionality Constraint on Supervenience: If P -properties supervene on Q -properties, then no two things can differ greatly with respect to their P -properties without differing greatly also with respect to their Q -properties.⁴⁶

As stated, this principle overgeneralizes. Let the Q properties be the family of microphysical properties. Let P be the singleton set containing the property of being an action that is not permissible. Consider two worlds, one where Darryl, fifty days in succession, diverts his attention from his child permissibly, but only just permissibly: in each case, if he had diverted his attention just one-hundredth of a nanosecond less, he would have acted impermissibly. The second world is extremely similar, microphysically, except that Darryl, fifty days in succession, diverts attention from his child impermissibly but only just impermissibly—on each day his attention is one one-hundredth of a nanosecond less than each corresponding day in the first scenario. Once we have made our peace with the sorites, classical propositional logic, and the relevant

46 Constantinescu, “Moral Vagueness,” 182.

inference rules for the quantifiers, it is hard to preclude pairs of possible worlds like this. By hypothesis, the two worlds do not differ much with respect to the *Q* properties. But they differ a good deal with respect to the *P* properties: there are fifty instances of *being an impermissible action* in one world but the corresponding actions in the other world do not instantiate that property. Similar issues come up with myriad choices of non-ethical *Ps*. Take the *P* properties to be the property of having a perfectly flat tabletop. Consider a pair of worlds, one in which there are fifty tables with perfectly flat tabletops, another with fifty tables that have vanishingly small imperfections in their tops so that none of them are perfectly flat. The underlying distribution of matter can be very similar but the difference in distribution of *P* properties is very significant. In sum, we cannot underwrite 3 by anything nearly as general as the principle that Constantinescu labels “Proportionality.” Nevertheless, there are some intriguing issues in the vicinity to which I now turn.

3. CONTINUITY

I articulated Constantinescu’s argument in terms of an inconsistent triad. Assuming some background logic of the sort previously alluded to, we are left to choose between a proportionality idea—namely, that (at a rough first pass) tiny physical differences in cases cannot render fitting significantly different attitudes and treatment—and the idea that the difference between permissible and not permissible actions (other things being equal) warrants markedly different treatments.

Something like the proportionality idea also gets advanced in Dougherty. Moved by Ted Sider’s thought that one “cannot both uphold epistemicism and continue to believe that differences in vague predicates always retain the significant we previously took them to have,” he writes:

There is some plausibility to thinking that if there is only a tiny descriptive difference between the actions, then any ethical difference could not be very important.⁴⁷

He concludes that, given epistemicism, we need to “scale back on the significance we place on applying these predicates.”⁴⁸ Here again it bears emphasis that the observation that tiny differences can make a moral difference do not rely on the full epistemicist package: some humble truths together with some standard inference rules for quantifiers and connectives all by themselves

47 Sider, “Hell and Vagueness,” 65; Dougherty, “Vague Value,” 361.

48 Dougherty, “Vague Value,” 362.

deliver the conclusion that tiny differences can separate the permissible from the impermissible. That said, an epistemicist commitment to semantic plasticity may in some cases play a distinctive motivating role in moving us to “scale back” our estimations of significance.

I shall propose a way of sharpening the proportionality idea that generates theses that can be explored with some rigor. (It will be a sharpening that does not really require the ideology of “tiny” and “significant.”)⁴⁹ The idea is to find ethical *scales* that seem of foundational ethical significance to actions, and to inquire as to whether, as one moves continuously from one physical possibility to another, the values on the ethical scale vary continuously. As a ridiculously simple but conceptually instructive toy model, pretend that the only way that reality varied across time and across worlds was on one parameter: the height of Jones. Take any path through time or worlds where Jones’s height varies continuously (so that, for example, there is no time in the series such that Jones is x inches tall at t but “jumps” in height so that at all times in some period after t Jones is, for some fraction $1/n$ of an inch, at least x plus $1/n$ inches tall). A proportionality thesis about some moral scale will say that the values on the ethical scale will vary continuously along that path as well. Now the world is obviously a lot more complicated than that. But we can still apply the same basic idea: once we have a notion of things varying continuously in a physical way along a path (without “jumps”), we have the means to state proportionality theses of the sort I am interested in. To fix ideas I shall be looking at some proportionality theses that focus on the microphysical terrain—here the notions of continuous variation have a natural home. A proportionality thesis will claim that the values on the ethical scale vary continuously as the microphysical terrain varies continuously. An anti-proportionality thesis will allow for discontinuous shifts. Let us look at a few proportionality theses (or perhaps, better still, “Continuity Theses”) along these lines

Some Moral Continuity Theses:

1. If a series of possible worlds vary continuously in their microphysical profile, then insofar as they vary in moral value, they will vary continuously in their moral value.⁵⁰

49 These ideas are touched on briefly in Dorr, Hawthorne, and Yli-Vakkuri, *The Bounds of Possibility*, 323. Discussions with Dorr have greatly influenced the writing of this section of the paper.

50 One might be tempted to instead articulate principles along the lines of, “If two cases have a small microphysical difference, then their difference on the moral scale is small,” and more generally, “The size of microphysical difference corresponds to the size of the difference on the moral scale.” (This kind of formulation is at least naturally suggested by Dougherty’s idea that “tiny descriptive differences” cannot make for a “very important”

2. If a series of possible people vary continuously in their microphysical profile, intrinsic and relational, then insofar as they vary with regard to their moral worth, they vary continuously with regard to their moral worth.
3. If a series of possible people vary continuously in their microphysical profile, intrinsic and relational, then, insofar as they vary with regard to how meaningful their lives are, they vary continuously with regard to how meaningful their lives are.
4. If a series of possible people vary continuously in their microphysical profile, intrinsic and relational, then (insofar as they vary at all) they vary continuously with regard to how much admiration is fitting for them. (There are obviously similar principles for other reactive attitudes.)
5. If the microphysical character of the world evolves continuously over a series of times, then insofar as the intrinsic moral value of each time varies, it varies continuously.

We should immediately acknowledge one way that these theses might naturally be weakened. As noted earlier, the classical magnets view is very tempting for phenomenal consciousness. It is, for example, very tempting (for atheists) to think that in the life of a person there is a metaphysically special time that marks the last time that the person is phenomenally conscious. Armed with this picture, someone might think that there is a very easy answer to such questions as, “How could a tiny microphysical shift mark a morally significant boundary?”—shifts that will seem insignificant under the lens of microphysics will mark metaphysically significant transitions from a case in which the very special

ethical difference, at least in a context where one is giving “descriptive” a microphysical spin.) But such principles are challenging for a number of reasons. For one thing, certain physical variations, even sizable ones (say to the configuration of sand dunes in an uninhabited desert), may make *no* moral difference whatsoever. One can fix for this by opting for the idea that the difference on the moral scale is *no more than* the difference on the physical scale (and indeed Dougherty’s own formulation is cautious in the required way). However, and perhaps more importantly, it is not clear how to give proper discipline to a notion of smallness that is cross-scale applicable. Claims like “*x* amount of time is as small as being *y* inches tall or as being *y* kilograms in mass” sound somewhat fishy, and if made at all rely on highly context-dependent, impressionistic reactions that are hard to systematize and risk relying on cross-scale notions that are too vague to be tractable. Of course there is no reason to think Dougherty is confining himself to microphysical facts when he says “descriptive facts.” But the issue raised above—how to generate cross-scale notions of smallness in a way that is not excessively vague and impressionistic—is still a somewhat pressing one. I leave the challenge of rigorously developing that kind of formulation to others. The continuity theses in the text do not rely on any such notions.

phenomenon of phenomenal consciousness is present to one where it is absent. Unless one wished to combat the classical magnets picture of consciousness, it seems that one will have to be more guarded when stating continuity. The more guarded version of continuity will say, of the relevant moral scale, that if a series of cases is continuous with regard to *both* microphysical and phenomenal consciousness, then the values along the relevant moral scale will, if they vary at all, vary continuously. I shall assume that insofar as readers are drawn to the relevant picture of phenomenal consciousness, they will operate with weakenings of this sort.⁵¹

For each pair of a property and a moral scale, there is an anti-continuity option, one that posits a discontinuous shift along the moral scale at certain points along a sequence of cases that vary continuously in microphysical respects. Consider, for example, a series of worlds involving Darryl that vary continuously in microphysical respects and where the key difference along the series is how long Darryl diverts his attention to his child. At one end of the series Darryl diverts his attention for a second, at the other for five minutes.

- 51 There are potentially far more guarded versions of proportionality theses. Consider a continuous path through worlds from one where I drive off this morning with a car that I own to a world where I drive off with a car that I do not own. One might think that, while the worlds continuously vary microphysically, there is still at some point a “jump” in the “descriptive facts”—from car ownership to non-car ownership. A far more guarded vision (for some moral scale) is along the following lines: when the descriptive facts (taken *en bloc*) vary smoothly—with no “jumps” along a path, the values of the moral scale vary continuously as well. (Here I am drawing on Dougherty’s preferred ideology of “descriptive facts.”) The key challenge here is to articulate what it is for the descriptive facts to vary smoothly. After all, on a standard conception of properties, for any pair of cases there will be infinitely many descriptive properties had by one but not the other. When should we count a case as a cliff/jumping point in the relevant sense? I opt for the more ambitious continuity theses in the text, in part because I can more easily see them being made rigorous (since the notion of continuous variation in the microphysical state of the world is a respectable one in physics) but also in part because they may surprisingly turn out to be true for various moral scales of importance (and so, for example, and surprisingly, it may turn out that in the car-ownership case, the fundamental moral scales do not jump discontinuously at the point where I drive off with a car that I do not own.) Still, I do not want to preclude our making sense of more guarded versions of proportionality than the one in the text: if the reader finds a way they are comfortable with to articulate some other version, I invite them to see how the themes of the text play out for that alternative version.

I should mention one further caveat. Suppose the moral scale is somewhat discrete, lacking the structure of the real line. Suppose for example that there are only ten thousand possible levels of admiration. There is still an analogue to the notion of varying continuously—namely, as one moves continuously across a series of physical cases, a shift from one level of admiration being fitting to another always proceeds via all the intervening levels of admiration. (It again bears emphasis that this thesis has absolutely no need for the ideology of “small” or “tiny.”)

Suppose we accept that there is a time t that marks the boundary between permissible and impermissible. Suppose for example, diverting attention for the period up to and including t is permissible, but that any longer period of diversion is impermissible. Then we might naturally think that as the physical landscape changes continuously, there is a discontinuous shift in some or all of the important moral scales—for example, perhaps there is a sharp drop-off right after the t -diversion world in the admiration fitting for Darryl. After all, in all the worlds in the series where the inattention was greater than the period up to and including t , he impermissibly diverted attention but in the other worlds he did not. According to the anti-continuity picture, a graph of the moral values across the series would display an abrupt cliff right at t . But the continuity lover for a particular moral scale will deny that the values from the scale are ever distributed in a cliff-like way across worlds that continuously vary in their microphysical makeup.

Other useful test cases for a continuity treatment of this or that moral scale are provided by properties that are not articulated using explicitly ethical language. Take the loving relation. One might think that if one possible individual loved no one while another possible individual loved someone, then, insofar as they were extremely close in moral worth, that would have to be because of some significantly compensating feature in the life of the loveless being—a feature that compensated for the absence of love. But reflection on sequences of worlds that vary continuously in their physical landscape, in combination with a continuity thesis about the moral-worth scale, suggests that this may be wrong. We can find a path from a lover to a non-lover that marches through continuously varying microphysical profiles. There will be pairs of cases separated by tiny microphysical differences, one of which involves the borderline presence of love, the other the borderline absence of love. And so there will be pairs arbitrarily close to each other in microphysical makeup that are divided by the absence and presence of love. If one subscribes to the continuity principle for moral worth, one will deny that moral worth abruptly drops off as one crosses the boundary to a loveless being.

Also consider principle 5, about the intrinsic value of times. Imagine a world containing a single creature who, at some point in time, dies. The defender of anti-continuity will naturally think that the intrinsic value of times sharply drops off at the point at which the creature is dead (especially if that was the only living creature left in the world).⁵² The defender of continuity for the

52 I take some liberties here and elsewhere—a more careful (but also more verbose) statement would acknowledge the difference between open and closed intervals. There may not be a first time where the creature is dead; rather, the period of life might form a closed interval and the period of death an adjacent open interval with no first time.

intrinsic value of times will say that there is no discontinuous jump in the intrinsic value of times.

The choice between continuity and anti-continuity for various moral scales is an intriguing one. The last example might initially seem to suggest that anti-continuity is obviously the way to go. Is it not obvious that at the very point someone dies there will be an abrupt jump in how much concern is due to them and so on, and if there only a few people in the world, an abrupt jump (downward) in the overall value of things? But on reflection things are far from obvious here. After all, it is overwhelmingly plausible on reflection that the line between life and death is vague. Supposing we take the semantic plasticity approach of the epistemicist, we will think that one cutoff is expressed, but myriad other cutoffs could very easily be expressed. It is moreover natural to think that candidate cutoffs trail off in terms of ease of being expressed by “the boundary between that creature’s life and death” in a continuous way. If none of the cutoffs are particularly natural or metaphysically privileged vis-à-vis their neighbors, then it seems plausible that moral value takes a continuous curve downward around the point of death rather than the cliff edge conducive to an anti-continuity picture.

The fan of proportionality will take a similar perspective on other cases: the boundary between loving and not loving may seem *prima facie* to enjoy such immense moral significance that there is an abrupt dip in value, fittingness of admiration, and so on, at the point in a physically continuous series of worlds that marks the boundary between a lover and a non-lover. But on reflection it seems plausible that “love” is vague, that there is no metaphysically privileged boundary in the vicinity, and on this basis one might well, for analogous reasons, begin to like the picture that the dip in value as love recedes across the series of worlds will, when graphically depicted, look like a curve rather than a sharp cliff edge.

Once one has gotten used to continuity in those cases, it is at least tempting to extend it even to the case of moral vocabulary—like “permissible”—that Dougherty focuses on. Consider the line between Darryl permissibly diverting attention and not permissibly diverting attention. Adopt a classical magnets approach and it is natural to say that there the line marks a discontinuous cliff in values along relevant moral scales. In that setting it is natural to combine an anti-continuity approach with the thesis that the exact place of a discontinuous jump is difficult to know. If someone presses, “How could a tiny physical difference mark a not-so-tiny shift in values along the relevant moral scale?” the classical magnets lover will insist that a tiny microphysical difference belies a metaphysically important difference that is hidden from view when one looks at the world merely through a microphysical lens. But once one takes an

epistemicist approach, acknowledging that the boundary between permissible and impermissible does not stand out as metaphysically significant vis-à-vis nearby boundaries, then it becomes far more tempting to think that the various moral scales—like moral worth and fittingness for admiration—do not have a cliff-like structure but rather have a curve-like structure around that boundary.

That said, I do not think the epistemicism is by any means *forced* to anti-continuity, an approach to which I now turn. Anti-continuity certainly has something to be said for it. It is far too easy to caricature the anti-continuity approach. Let us return to Constantinescu's thought that it is "unjust to praise one person and blame another when the difference between their actions was slight." It may be thought that an anti-continuity approach to the reactive attitudes would recommend heaping praise on Darryl when he just about manages to act permissibly, and blame, contempt, censure, and so on when he just about fails to act permissibly. (Indeed Constantinescu's talk of blaming and chastising the person who only just about acts impermissibly encourages this vision.) But even if anti-continuity is right, that is the wrong picture. After all, when Darryl acts only just permissibly, he acts in a way that is, for all he knows, impermissible. That is not great. And when Darryl acts only just impermissibly, he does not know he is acting impermissibly. That is not nearly as bad as knowingly acting impermissibly. Moreover, it is natural to think that norms such as "Don't divert attention impermissibly" generate what Williamson calls "derivative norms," including, in this case, the "secondary norm" that people should have the disposition not to divert their attention impermissibly.⁵³ Darryl's only just permissibly diverting his attention may be a tell-tale sign that he does not have a stable disposition to permissibly divert his attention and thus signal failure to comply with the secondary norm. Suppose instead that coming close to acting impermissibly was an aberration and that Darryl does have a stable disposition to act permissibly (that was interfered with in an unusual way on this occasion). Then the relevant case of Darryl's acting impermissibly, since it is ever so similar, will likely also be an aberration, and in that case, too, the secondary norm will be satisfied.

All this suggests that we should not think of the difference between the two cases as all that great. We should not think it is a boundary so dramatic as to warrant something as contrasting as heaping of praise versus heaping of blame.⁵⁴ But such a concession is perfectly compatible with the claim that the boundary between the permissible and the impermissible marks a discontin-

53 Williamson, "Acting on Knowledge."

54 Of course, this kind of response is not available to a theologian who thought that each possible life warranted one of two sharply dichotomous divine reactions, being sent to heaven and being sent to hell. Thus, even the proponent of anti-continuity might not be very comfortable with the thought that cases almost the same physically can warrant

uous boundary with regard to the fittingness of various reactive attitudes. Not knowing where the boundary lies, the lover of anti-continuity will then be committed to not knowing exactly what degree of this or that reactive attitude is most fitting in certain cases. But that is all right. Having made their peace with unknowable moral truths, it is not a particularly great additional shock to make one's peace with the fact that among such truths are facts about exactly which levels of which reactive attitudes are fitting in various cases.

If anti-continuity is defensible for this or that moral scale in connection with a property like being permissible (so that the cliff edge on the scale lines up with the boundary of the property), the key to defending such a view will, I think, be to promote thoughts such as the following: the whole point of various moral predicates is they warrant at least somewhat significant differences in reactive attitudes, mark somewhat significant differences in moral worth, and so on. One can run sorites series on such predicates as "being evil" and find physically continuous series where there is a cutoff associated with these predicates. But if the whole point of these predicates is that they warrant at least somewhat significant differences in reactive attitudes, and trigger somewhat significant differences with respect to other important moral scales, then their impact on certain of those scales had better be cliff-like.

How is this thought to be reconciled with a semantic plasticity claim to the effect that it is a semantically fragile matter which line is drawn by "evil," "permissible," and so on? The natural way to harmonize things here is to posit a kind of penumbral connection between the relevant predicates and the language of the scale.⁵⁵ Suppose, for example, in keeping with anti-proportionality, one held that the boundary on a sorites series for "being an evil person" also marked a discontinuity in the scale associated with the question, "How fitting it is to hate that person?" And suppose, in keeping with semantic plasticity, that one thought that at nearby worlds, "being an evil person" expresses different properties that draw a boundary on the same sorites series in slightly different places. A natural thought that the meaning of the question, "How fitting it is to hate that person?" also shifts in a way that is coordinate with the shift in "evil person." At the actual world, the scale associated with "How fitting it is to hate?" expresses a scale that shifts discontinuously around boundaries that mark the difference between evil and not being evil (of course the discontinuous shift may be not as dramatic as

treatments as radically different as being sent to Heaven or Hell. This issue is explored at length in Sider, "Hell and Vagueness."

55 "Penumbral connection" is the expression typically used by supervaluationists and epistemicists to mark logical, analytic, or *a priori* connections between predicates that are respected by all families of candidate interpretations. Penumbral connections obtain when individually admissible interpretations for several words are not jointly admissible.

one initially imagines because the first evil people in the series are only just evil, and the last are almost evil). Meanwhile at a nearby world where “evil” means some slightly different property, the scale associated with “How fitting is it to hate?” is different to the actual scale, by in particular having a little cliff that marks the boundary associated with that slightly different property (and thus, presumably, there is a slight shift in the meaning of “fitting” that coordinates with the slight shift in the meaning of “evil”). This view of the matter is not obviously the right one. But I submit that it is not obviously the wrong one either.

That said, I would take quite a bit of persuading in many cases to opt for anti-proportionality. Let me illustrate the relevant issues by looking at a few test cases.

Imagine someone that badly wanted a fantastic kitchen. They begin with a humdrum kitchen that gradually improves over the years, and their attitude toward their kitchen slowly evolves. At some point it is clearly fantastic. The surfaces had begun to sparkle more and more and there came a point where, given all the other changes, it just so happened that a tiny bit of extra sparkle took it over the edge to being fantastic, after which it continued to improve further. Looking back they see that there were no discontinuous shifts in their attitude to their kitchen, though by the end their attitude is extremely positive and at the beginning very negative. What should we say if, looking back, the person said to themselves, “I didn’t really notice the point that marked the boundary between the kitchen not being fantastic and being fantastic. What was fitting at that point was an extra little celebration and hence a discontinuous jump in positivity. After all, what I always wanted was a fantastic kitchen!”? Someone making the speech might think to themselves, “Granted, my attitude toward my kitchen did not in fact take a noticeable jump at the point where it became fantastic. But that is because I did not know when that shift occurred. While there is no particular point in the renovation process that I can point to as a point where the shift from non-fantastic to fantastic occurred, I nevertheless know that there *was* a point where this shift did occur. And whenever it did, a marked jump in positivity of attitude was *fitting*.”

A self-acknowledged desire for a fantastic kitchen is a pretty good rough and ready way to frame one’s domestic ambitions. Yet I am dubious that there are good grounds for anti-proportionality here. The natural way to develop anti-proportionality for the scale of, say, fitting pride, is to suppose a penumbral connection between the question, “How much pride is fitting?” and the meaning of “fantastic,” so that as the cutoff associated with the word “fantastic” moves around in nearby worlds, there is also a corresponding shift in the meaning of the question, “What level of pride in your kitchen is fitting?” (After all, it is hard to imagine that the shift from non-fantastic to fantastic would

be significant for the scale associated with the expression “the level of pride that is fitting” at a world where “fantastic” meant some other property, fantastic*.) Here, then, is the anti-proportionality vision: given the actual meaning of “the level of pride that is fitting,” a certain discontinuous leap in level of pride in one’s kitchen is fitting as the kitchen moves from not-fantastic to fantastic. Meanwhile, the meaning of “the level of pride that is fitting” slightly shifts its meaning at those nearby worlds, where “fantastic” slightly shifts its meaning in such a way that the meaning of “fitting pride” at *those* worlds calls for no such jump at the point at which the kitchen transitions from non-fantastic to fantastic. But, for what it’s worth, the hypothesized penumbral connection does not seem to me to be especially plausible. The meaning of the question, “How much pride in your kitchen is fitting?” does not seem hostage to the semantic vicissitudes of “fantastic” in this way.

A second test case: someone slowly descends into depravity over their lifetime so that at some point they become evil and their parents become slowly more disgusted and ashamed of them. What should we say if the parents, looking back make the speech, “We didn’t really notice the exact point that marked the boundary between our child not being evil and being evil. But at whatever point that shift occurred, a discontinuous jump in negative attitudes, including a marked extra degree of moral disgust, was warranted”? On reflection the case does not seem so different from the case of the kitchen. If you agree with me that it is a bit silly to posit a penumbral connection between “How much pride is fitting?” and “fantastic” that requires a discontinuous shift in what attitudes are fitting once the boundary for being fantastic is crossed, it is arguable that it is similarly implausible to posit a penumbral connection of this sort in the case of being evil.

I have voiced some hesitation about anti-continuity ideas, expressing considerable sympathy with Dougherty’s idea—following Sider—that an epistemicist treatment of vague predicates “may involve scaling back the significance we place on applying these predicates.”⁵⁶ However, I have provided nothing like a knockdown argument. Indeed, I do not wish to be doctrinaire here. The issues are certainly very delicate. Nor do I wish to assume a monolithic approach to continuity. For any pair of a property that can divide a physically continuous series of cases and a moral scale, one can ask whether the moral scale varies discontinuously around the boundary marked by the property or not. Perhaps a systematic defense of continuity theses for the important moral scales is possible here, where considerations of vagueness will in many cases play an important role in the defense. But, as I have indicated, anti-continuity is not

⁵⁶ Dougherty, “Vague Value,” 11.

dead in the water, and perhaps a selective defense of anti-continuity for some property/morally significant scale pairs is defensible. My aim here has not been to resolve this matter; it is to bring to attention a helpful way of thinking about proportionality issues that grounds them in questions of continuous variation. This approach to proportionality is one that, as far as I can tell, has not found that much life in the ethics literature thus far.

4. CONCLUDING REMARKS

I have found no good reason as yet to think there is anything deeply problematic about an epistemicist treatment of moral predicates. There are, to be sure, arguments in the literature to the effect that such an approach to vagueness in ethics is problematic. But these arguments are wanting. Many such arguments, if they work at all, work against any approach to vagueness that assumes classical propositional logic and some standard rules for the quantifiers, of which epistemicism is but a species of a wide genus. My hunch is that the relevant authors have not made their peace with a vital choice point: reject some mundane inference rules or else simply accept that, even when it comes to moral properties such as permissibility, tiny differences can make the difference between instantiation or non-instantiation. Nevertheless, a number of the critical discussions of epistemicism about moral vagueness point us toward hugely interesting choice points in various subdomains of ethical theorizing, ones that turn on whether to think that, as cases vary in a physically continuous ways, the boundary associated with some property marks a discontinuous shift in the values along some moral scale. Once we have made our peace with classical propositional logic and some orthodox inferences rules for the quantifiers, worries about unknowable truths come to nothing, as do worries that a tiny physical difference cannot make any moral difference. But these proportionality questions remain and the question of how to resolve them is both pressing and intriguing. I have indicated how an epistemicist approach might begin to approach them. And I would encourage those readers who opt for some alternative account of vagueness to explore proportionality issues within the framework of their favored approach.⁵⁷

University of Southern California and Australian Catholic University
jhawthor@usc.edu

57 I am grateful to Cian Dorr, Stephen Finlay, Miriam Schoenfield, two anonymous referees, and a discussion group at Australian Catholic University for helpful comments and discussion.

REFERENCES

- Bacon, Andrew. *Vagueness and Thought*. Oxford: Oxford University Press, 2018.
- Constantinescu, Cristian. "Moral Vagueness: A Dilemma for Non-Naturalism." *Oxford Studies in Metaethics*, vol. 9, edited by Russ Shafer-Landau, 152–85. Oxford: Oxford University Press, 2014.
- Davies, Martin, and Lloyd Humberstone. "Two Notions of Necessity." *Philosophical Studies* 38, no. 1 (July 1980): 1–30.
- Dietz, Christina. "Reasons and Factive Emotions." *Philosophical Studies* 175, no. 7 (July 2018): 1681–91.
- Dorr, Cian, John Hawthorne, and Juhani Yli-Vakkuri. *The Bounds of Possibility: Puzzles of Modal Variation*. Oxford: Oxford University Press, 2021.
- Dougherty, Tom. "Vague Value." *Philosophy and Phenomenological Research* 89, no. 2 (September 2014): 352–72.
- Dunaway, Billy, and Tristram MacPherson. "Reference Magnetism as a Solution to the Moral Twin Earth Problem." *Ergo* 3, no. 25 (2016): 639–79.
- Finlay, Stephen. *Confusion of Tongues: A Theory of Normative Language*. Oxford: Oxford University Press, 2014.
- Fine, Kit. "The Possibility of Vagueness." *Synthese* 194, no. 10 (October 2017): 3699–725.
- . "Vagueness, Truth and Logic." *Synthese* 30, nos. 3–4 (April–May 1975): 265–300.
- Fritz, Peter, John Hawthorne, and Juhani Yli-Vakkuri. "Operator Arguments Revisited." *Philosophical Studies* 176, no. 11 (November 2019): 2933–59.
- Grice, Paul. *Aspects of Reason*. Oxford: Clarendon Press, 2001.
- Hawthorne, John. "Epistemicism and Semantic Plasticity." In *Metaphysical Essays*, 185–210. Oxford: Oxford University Press, 2006.
- Hawthorne John, and Ofra Magidor. "Reflections on the Ideology of Reasons." In *The Oxford Handbook of Reasons and Normativity*, edited by Daniel Star, 113–42. Oxford: Oxford University Press, 2018.
- Kearns, Stephen, and Ofra Magidor. "Epistemicism about Vagueness and Meta-Linguistic Safety." *Philosophical Perspectives* 22 (2008): 277–304.
- McGrath, Sarah. "Moral Realism without Convergence." *Philosophical Topics* 38, no. 2 (Fall 2010): 59–90.
- Parfit, Derek. *On What Matters*, vol. 1. Oxford: Oxford University Press, 2011.
- Schoenfield, Miriam. "Moral Vagueness Is Ontic Vagueness." *Ethics* 126, no. 2 (January 2016): 257–82.
- Sider, Ted. "Hell and Vagueness." *Faith and Philosophy* 19, no. 1 (January 2002): 58–69.
- Sorensen, Roy. *Vagueness and Contradiction*. Oxford: Oxford University Press, 2001.

- Unger, Peter. *Ignorance: A Case for Scepticism*. Oxford: Oxford University Press, 1979.
- Williamson, Timothy. "Acting on Knowledge." In *Knowledge First: Approaches in Philosophy and Mind*, edited by J. Adam Carter, Emma C. Gordon, and Benjamin Jarvis, 163–81. Oxford: Oxford University Press, 2017.
- Williamson, Timothy. *Vagueness*. Milton Park, UK: Routledge, 1994.
- Wright, Crispin. "Truth in Ethics." *Ratio* 8, no. 3 (December 1995): 209–26.
- Yli-Vakkuri, Juhani. "Epistemicism and Modality." *Canadian Journal of Philosophy* 46, nos. 4–5 (August 2016): 803–35.