

Moral Vagueness, Epistemicism, and The Problem of Infectious Vagueness

Ofra Magidor

(forthcoming in *Lauener Library of Analytical Philosophy*)

Abstract:

It has been suggested that the following three theses are incompatible: Moral Realism, Epistemicism about vagueness, and the claim that moral terms are vague. If this is so, (at least) one of these three must be rejected. This paper explores the possibility of resolving this trilemma by rejecting Moral Vagueness.

Such a rejection is no easy feat because vagueness is seemingly infectious: in several cases, an uncontroversially vague descriptive claim is (definitely) materially equivalent to a moral claim, and thus the vagueness for the descriptive terms threatens to render the moral terms vague as well.

However, I argue that Epistemicism has the resources to respond to this problem of infectious vagueness. In previous work, I argued that epistemicists should reject a key principle in the logic of vagueness: Distribution, which states that if a biconditional is definitely true, then if one side of it is borderline, then so is this other. Here, I expand that argument to show how epistemicists can coherently reject not only the principle in general, but the particular instances of it that link moral and descriptive statements. This allows epistemicists to maintain that even when the relevant descriptive statement is borderline, the materially equivalent moral statement can nevertheless be definite and non-vague, thus allowing the epistemicists to reject Moral Vagueness and uphold Moral Realism.

Introduction

Several arguments in the literature purport to establish the following claim:¹

Incompatibility: the following three theses are jointly incompatible:

(i) Moral Realism (ii) Epistemicism (iii) Moral Vagueness.

Let us elucidate the relevant terms, starting with ‘Moral Realism’. Let *minimal moral realism*, be the thesis that moral statements are truth-evaluable, and in many cases are true. Let *moderate moral realism* be the conjunction of minimal moral realism with the claim that the truth-makers for moral statements are (at least in the majority of cases) independent of human judgements or reactions.² Let *robust moral realism* be the conjunction of moderate moral realism with the additional claim that “ethical properties are part of the deep underlying metaphysical structure of the world” (Dougherty (2014, 358)). Since some of the arguments for Incompatibility (including the one I favour) rely on this strong interpretation, I will take ‘Moral Realism’ to mean *robust moral realism*. Having said that, robust

¹ The most explicit argument for Incompatibility is due to Schoenfield (2016), but see the discussion below for a range of other arguments that point towards a similar conclusion.

Schoenfield actually argues for a stronger claim: that Moral Realism and Moral Vagueness are incompatible with any non-ontic view of vagueness (including epistemicism), and suggests that the best way to avoid this trilemma is to opt for an ontic view of vagueness. One upshot of my argument is that it offers a different way to avoid Schoenfield’s trilemma, one not requiring an ontic view of vagueness.

² The caveat in parenthesis is due to cases like this: all else being equal, I ought to offer you the ice-cream flavour you prefer, rather than the ice-cream flavour you dislike. But this fact depends in part on human judgements and reactions (i.e. your preference for this particular flavour). Clearly, this type of case is no threat to moderate moral realism.

moral realism only plays a role in motivating the discussion (§1). The remainder of the paper (§2-4) only presupposes minimal moral realism.

Next, *Epistemicism* is an approach to vagueness according to which even vague words such as ‘tall’ or ‘bald’ have sharp semantic-values, and sentences such as ‘Jill is tall’ or ‘Harry is bald’ conform to the laws of excluded middle and of bivalence even when they constitute borderline cases. Moreover, in this paper I will assume a variant of epistemicism which broadly follows the classic development of the view by Williamson (1994). According to this variant of epistemicism, vague terms are *semantically plastic*: small shifts in their use make for small shifts in their meaning. Furthermore, the semantic plasticity of vague terms explains why we do not and cannot know whether they apply in borderline cases. (More details on how this explanation of ignorance works in §4.1 below.)

Finally, *Moral Vagueness* is the claim that the basic pure moral terms are vague. For concreteness, I will assume that these include the terms ‘ought’, ‘permissible’, and ‘required’.³ The restriction to ‘basic pure’ moral terms is important: even those who maintain that terms such as ‘ought’ or ‘permissible’ aren’t vague would presumably agree that some more derivative or hybrid moral terms are vague (consider for example ‘somewhat permissible’, or thick moral terms such as ‘evil’). For ease of readability, I drop this qualification in the remainder of the paper.

The paper is structured as follows. In §1, I briefly examine several arguments raised in extant literature in support of Incompatibility. I argue that most of these arguments are found lacking, but there is nevertheless one compelling motivation that supports Incompatibility. This suggests that (at least) one of the three theses ought to be rejected. The remainder of the paper explores one way of addressing the trilemma: a view that upholds Epistemicism and Moral Realism, while rejecting Moral Vagueness.

In §2, I explain why even though the existence of Moral Vagueness might seem intuitive, rejecting it is hardly a non-starter.

However, in §3 I present a much more serious challenge for the rejection of Moral Vagueness: *the problem of infectious vagueness*. The worry is that moral statements are sometimes (definitely) materially equivalent to descriptive statements, ones which are themselves uncontroversially vague. For example, suppose I promise to give my friend ten pounds if they date a bald person. Following the promise, I am required to give my friend ten pounds if and only their date is bald. The worry is that the vagueness of the descriptive statement (‘The date is bald’) infects the moral statement (‘I am required to give my friend ten pounds’), and thus the moral statement (and in particular the moral term ‘required’), must be vague as well.

In §4, I show how epistemicists can respond to the problem of infectious vagueness. In previous work (Magidor 2018), I argued that epistemicists should reject a key principle in the logic of vagueness: Distribution, which states that if a biconditional is definitely true, then if one side of it is borderline, then so is this other. Here, I expand that argument to show how epistemicists can coherently reject not only the principle in general, but the particular instances of Distribution which apply to biconditionals linking moral and descriptive statements. This allows epistemicists to maintain that even when the relevant descriptive statement is borderline, the materially equivalent moral statement can nevertheless be definite and non-vague, thus allowing the epistemicists to reject Moral Vagueness.

³ Of course, in addition to their potential vagueness, these terms are most likely context-sensitive. I fix here on an interpretation of them that involves *moral* (rather, than say, prudential) requirements and takes them to have the ‘all things considered’ interpretation.

§1 Motivating the incompatibility of Epistemicism, Moral Realism, and Moral Vagueness

The literature contains a range of arguments in favour of Incompatibility. The crux of the issue is the assumption that, according to Epistemicism, if moral terms are vague then they must be semantically plastic: the alleged incompatibility is really one between Moral Realism and the semantic plasticity of moral terms.⁴

Why should the claim that moral terms are semantically plastic seem troublesome? Such plasticity entails that there is a possible community of speakers who use these terms only slightly differently from the way we do, and where the terms in their dialect pick out slightly different properties. Thus, for example, it might turn out that when we utter the sentence ‘One is required to give n percent of one’s income to charity’ we speak truly, but when our counterparts utter the same sentence, they speak falsely - merely because ‘required’ in their mouths picks out a slightly different property (call it ‘required*’).

Several arguments have been raised for thinking this kind of linguistic variance is troubling, especially against the background of Moral Realism. Few of these arguments, however, ultimately survive scrutiny.

One worry maintains that given such variance, any dispute with members of this alternative community in which we state that some action is ‘permissible’ and they state that it is ‘not permissible’ would be a merely verbal dispute and effectively pointless: both sides would be simply talking past each other. This conflicts with the strong intuition that such arguments can often reflect substantive disagreement about serious moral questions.⁵

But it is not clear that our dispute with members of the alternative community would be ‘merely verbal’, and it certainly isn’t obvious that it would be “rather silly” (Horgan & Timmons 1991, 460). The question of what counts as a merely verbal dispute is itself delicate and beyond the scope of the current paper.⁶ But it is worth keeping in mind that two speakers can presumably have a substantive argument even if they use terms with slightly different meanings. For example, suppose we’re arguing about whether ‘John is tall’. We’ve both only seen John from a distance. I estimate that he is roughly 1.6m in height, while you estimate that he’s 1.9m in height. Even if the cut-off for ‘tall’ is 1.81 in my

⁴ It’s worth noting that one way to develop Epistemicism maintains that while paradigmatic cases of ignorance due to vagueness are explained via semantic plasticity, there are some more derivative forms of vagueness that do not require plasticity (see e.g. Williamson 2016, §3 and Manley *ms.*, 12).

Adopting this alternative, does not affect the gist of my arguments in this paper, though some modifications to its framing would be needed.

As noted above, the worries raised in this section ultimately aim to establish the incompatibility of Moral Realism and the semantic plasticity of moral terms. Moreover, the problem of infectious vagueness (§3) can be reframed as the problem of *infectious semantic plasticity*. The arguments in this paper can then be taken to show that the semantic plasticity of descriptive terms does not infect the plasticity of moral terms, and hence does not pose a problem for Epistemicism. (See also n. 35 below on further remarks on how to formulate such a reframing.)

⁵ See Schoenfield (2016, 268), Dougherty (2017, 189), and Manley (*ms.*, 6-9). The worry can be traced back to Horgan & Timmons (1991), who raise it not in the context of vagueness, but rather of other views on which moral terms might differ slightly in meaning in nearby environments. See also the discussion in Dunaway & McPherson (2016, §4).

⁶ See Balcerak Jackson (2020) for an overview of the literature on verbal disputes.

dialect, and 1.82 in yours, we could presumably still be having a substantive disagreement about this matter.⁷

A second worry, goes like this:⁸ what makes a term such as ‘permissible’ have different referents in the two linguistic communities are the different use facts of the two communities. Moreover, a particularly central role in the determination of the meaning of these terms is arguably played by use facts of sentences of the form ‘x is/isn’t permissible’, including of course which sentences of this form members of each community accept or reject. But then, the suggestion goes, the current view conflicts with robust moral realism. This is because it the very fact that members of a community have certain attitudes towards what counts as ‘permissible’ plays a crucial part in ensuring that such assertions are true. This seems to suggest at best a kind of ‘stance dependent’ rather than robust moral realism.

But as Manley (*ms.*,6) correctly points out, this worry rests on a conflation between semantics and metasemantics. The fact that speakers’ attitudes might end up (in part) determining which properties are picked out by their moral terms, is part of the metasemantic question of how moral terms end up getting the meanings they have. But this says nothing at all about the truth-makers for the moral statements themselves, which is the question at stake for Moral Realism.

A third worry is that those cases where we face difficult moral questions are arguably going to align precisely with borderline cases of vague moral terms. But then it might seem hard to make sense of our practices of moral deliberation. Suppose, for example, we are agonising over the question of whether it is permissible to abort a foetus at a certain stage of development. It might seem that according to the epistemicist, there is no need to agonise: the question just turns on a fairly superficial matter of which property out of several nearby candidates our word ‘permissible’ picks out, and moreover given that what determines the answer to this question are use facts, it might seem that the moral question can be resolved simply by consulting a linguist and “crunching through linguistic data” (Schoenfield 2016, 265).⁹

But, *pace* Schoenfield, it is not true that on the epistemicist view the question of what is permissible can be settled merely by surveying linguistic use facts: a key assumption of the framework is that part of our ignorance of the sharp cut-offs of vague terms includes ignorance of the way in which meaning facts depend on use facts, so even the most comprehensive survey of use facts would not be sufficient to settle the question.¹⁰

Finally, one might worry that the semantic plasticity of moral terms makes it hard to explain why we take morality to play crucial roles such as being action guiding: it seems that our actions could have just as easily been guided by what’s permissible* or required* rather than what is permissible or required.¹¹

⁷ A similar observation that so called verbal disputes can still be substantive is argued for in Balcerak Jackson (2020: 124-126).

⁸ See the discussion in Dougherty (2014, 360) and Manely (*ms.* 5-6).

⁹ This worry is also mentioned in Dougherty (2017, 189) and Liao (*ms.*, 13).

¹⁰ See Abasnezhad (2023, §2.1) and Hawthorne (2022, 224-25).

¹¹ A different worry which is sometimes raised against an epistemicist account of moral vagueness, (one not specifically related to semantic plasticity), is that it entails that there can be unknowable moral truths, and hence morality cannot be action guiding (see Wasserman 2012, 70; Dougherty 2014, 361; and Constantinescu 2014, 176). I do not find this worry particularly compelling either. For one thing, nearly every view of moral vagueness entails that there are situations where it is unknowable what one ought to or is permitted to do (even if there are technically no ‘truths’ involved): after all, on nearly any account of vagueness, borderline statements ‘S’ are ones where one cannot know that S and cannot know that not S. Moreover, as Liao (*ms.*,18) points out, the idea

Yet while it is tempting to think that permissibility is special because it (and not permissibility*) plays the role of being action guiding, we need to keep in mind that if moral terms are semantic plastic, then it is highly plausible that so are a range of related terms such as ‘action guiding’. After all, when we say that moral truths are ‘action guiding’ we presumably interpret this as a normative rather than a descriptive claim, meaning something like ‘we ought to act according to what morality dictates’. But then when our counterparts in the alternative community utter the claim ‘moral truths are action guiding’ they are also expressing a truth: that one ought* to act according to what is permissible*. So, while permissibility might be unique in playing the action guiding role, permissibility* might be unique in playing the action-guiding* role.¹²

So far, the case for Incompatibility looks shaky. Nonetheless, there is one remaining line of argument which, in my view, gives at least a strong motivation to think that, given (robust) Moral Realism, moral terms are not semantically plastic.¹³

Suppose we endorse a theory of reference along the lines of Lewis’s theory of eligibility.¹⁴ According to this view, there are two factors that go into determining the reference of a term: *fit* – i.e. how well a potential referent fits with our use facts, and *eligibility* – i.e. is how natural or metaphysically joint-carving a potential referent is. In particular, highly eligible properties can act as *reference magnets*: properties that are eligible enough to trump various discrepancies in fit, and hence are likely to be the stable referents across a wider variety of different uses. For the epistemicist, this theory can nicely explain the difference between vague terms such as ‘bald’ and ‘tall’ and precise terms such as ‘prime’ or ‘electron’:¹⁵ since there are a variety of equally eligible height properties that are potential referents for ‘tall’, its referent is very sensitive to use facts and is thus semantically plastic; by contrast, since there are unique most-eligible candidate referents for ‘electron’ or ‘prime’, these terms are not plastic.

Assuming we accept this theory of reference, where should moral terms fall? Given robust moral realism, they more plausibly fall into the latter category: moral terms pick out metaphysically special properties, which in turn suggests they act as reference magnets. As Wasserman (2012, 78) puts it: “For the serious moral realist, moral terms pick out metaphysically important properties. The property of being morally obligatory, for example, is more like the property of being gold, than being bald. So, the semantics of moral terms should be more like the semantics of natural kind terms. So, moral terms are not semantically plastic.”

There is thus at least one strong motivation for accepting Incompatibility, which means that one of (i)-(iii) must be rejected. The remainder of the paper explores one response to this trilemma: accepting Epistemicism and Moral Realism, while rejecting Moral Vagueness.

§2. Why rejecting Moral Vagueness isn’t a non-starter

that only knowable truths can be action guiding requires a highly internalist conception of the latter, one that in the light of anti-transparency considerations argued for in Williamson (2000) might not be feasible. See also Hawthorne (2022, §2.2) for further discussion of this issue.

¹² See Eklund (2017) for an extensive discussion of related issues.

¹³ See Wasserman (2012) and Dunaway & McPherson (2016) for sympathy to the idea that moral terms form reference magnets (though interestingly, in Dunaway & McPherson 2016, n. 15 they suggest that an exception might be needed specifically to accommodate moral vagueness).

¹⁴ Lewis (1983). Williams (2018) argues for a much more nuanced interpretation of the Lewisian metasemantic picture, but one on which moral terms are still not plastic.

¹⁵ Cf. Williamson (1994, 231).

At a first pass, it might seem that rejecting Moral Vagueness is a non-starter. Cases like the following might seem to straightforwardly demonstrate that moral terms are uncontroversially vague:

Incubator: A futuristic incubator contains a human sperm and an egg. If no one interferes, the incubator will initiate conception and care for the developing organism for nine months, until it develops into a human child similar to a typical baby at birth. Until what point in the process exactly is it permissible to destroy the contents of incubator?¹⁶

Charity: Jill has an option of donating some of her money in order to save the life of her colleague, a life that no one else is in a position to save. What amount of money exactly is such that, if that amount is needed to save the colleague's life, Jill is required to donate for this purpose?

Promise: Jack promises to give his friend ten pounds if the next person they go on a date with is bald. Exactly how many hairs does the date need to have in order for Jack to be required to pay his friend?¹⁷

It is certainly tempting to treat all of these cases as instances of vagueness. Two informal characteristics of vague terms is that they give rise to borderline cases and that they give rise to sorites arguments.¹⁸ Consider the case of 'tall'. It gives rise to borderline cases: there are people of intermediate height, where (even knowing the background facts such as their precise height) we are not sure whether they are tall or not tall and do not even know how we might go about settling this question. And the term is also susceptible to sorites arguments: someone who is one meter in height is clearly not tall, someone who is two meters in height is clearly tall, and yet the tolerance premise - 'if someone who is n meters in height is not tall, then someone who is n meters and one nanometre in height is not tall' - seems at a first-pass attractive (even if on reflection, we realise it ought to be rejected).

Both of these characteristics might seem to apply to moral cases as well. First, moral terms seem to give rise to borderline cases. For example, in Incubator, there seem to be intermediate points in the development process of the organism where (even if we know all the background biological facts), we are not sure if it is permissible to destroy it, and where we do not even know how might go about settling this question. And second, moral terms seem susceptible to sorites arguments. For example, in Charity, it seems clear that Jill is required to give one penny, clear that she is not required to give her entire income, and the relevant tolerance premise - 'if Jill is required to give n pounds to save the life, then she is required to give n pounds and one penny' - seems initially attractive.

However, these considerations in favour of taking moral terms to be vague are hardly decisive. Telling apart which cases of uncertainty or ignorance are due to vagueness and which are due to other sources is no easy matter. Moreover, the two characterisations mentioned above should not be treated as definitive sufficient conditions for the presence of vagueness (at least not in the rough form presented above). There are cases of terms that at least superficially seem to fit these characteristics,

¹⁶ This case is taken from Manley (*ms.*). There are obviously non-fictional analogues of this example that involve pregnancy and abortion, but I prefer this version of the case as it abstracts away from some of the moral complications of having a pregnant person.

¹⁷ One might feel that this case is different from the others, because here it is the vagueness of 'bald' rather than of 'required' that is at stake. But the worry – which is at the centre of this paper – is precisely that 'required' might inherit this vagueness from 'bald'.

¹⁸ See Eklund (2007).

despite not being vague. For example, consider the term ‘largest twin prime’:¹⁹ there is an intermediate case (the largest prime discovered thus far) for which we are not sure whether the term applies or even how we might go about settling the question.²⁰ But the term is clearly not vague – it is simply defined using mathematically precise language. Similar observations can be said about apparent sorites susceptibility. Suppose for example we are piling some sand on a bridge. Consider the complex term ‘sufficient amount of sand to break the bridge’. The relevant tolerance principle - ‘If n grains of sand aren’t sufficient to break the bridge then $n+1$ grains aren’t sufficient to break the bridge’ - does seem initially attractive.²¹ But arguably, the term is not vague: there is no relevant semantic indecision or plasticity in the vicinity, merely a precise yet complicated physics question about what weight exactly – and correspondingly how many grains - the bridge can sustain before breaking.²²

That is not to say that there couldn’t be more sophisticated characterisations of vagueness (or more precise delineations of the above characterizations) that would tell apart cases of vagueness from other types of ignorance. The point is merely that noticing some superficial similarities between paradigmatic cases of vagueness and cases such as Incubator, Promise, or Charity isn’t sufficient to conclusively establish the presence of moral vagueness. Rejecting Moral Vagueness is a serious option, and given the considerations raised in §1 above, an option well worth exploring.

§3. The problem of infectious vagueness

Unfortunately, the rejection of Moral Vagueness faces a much deeper challenge: *the problem of infectious vagueness*. There are cases where a moral statement²³ is materially equivalent to a descriptive statement, one which is itself uncontroversially vague.²⁴

In order to make our discussion concrete, consider three biconditionals which, given the relevant background moral commitments, one might wish to accept:²⁵

(Biconditional-1): In Incubator, it is permissible to destroy the entity in the incubator if and only if it is not a person.

¹⁹ A twin prime is a prime number n such that at least one of $n+2$ and $n-2$ is also prime (for example, 41 and 43 are twin primes). It is an open mathematical question whether there are infinitely many twin primes, and hence which number, if any, satisfied this predicate.

²⁰ Of course, we could settle the question in one direction if we discover a larger twin prime. But assuming we never manage to find such a prime or to prove that there are infinitely many twin primes, then it is not clear how we would settle the question.

²¹ Of course, further reflection shows that the tolerance principle has to be false, but the same is true of sorites arguments involving genuinely vague terms, especially in the context where we are accepting epistemicism.

²² Admittedly, there might be some lingering vagueness about counts as the bridge ‘breaking’. But the point is that even if we fully precisify this notion, there is something initially attractive about the tolerance principle.

²³ By ‘moral statement’, I simply mean a statement involving a moral term. However, as will be clear from the ensuing discussion, the crucial question is whether the moral statement is vague specifically due to the vagueness of the moral terms featuring in it.

²⁴ The problem of infectious vagueness isn’t often explicitly mentioned in the literature, but informal remarks regarding why we should believe in moral vagueness often appeal to the vagueness of the relevant descriptive conditions. One exception is Wasserman (2012) who explicitly argues that if personal identity is indeterminate then so are some moral properties.

²⁵ Of course, not everyone will accept these particular biconditionals: accepting them depends on one’s specific moral commitments. But it is easy to generate other variant biconditionals of this form, depending on one’s first-order moral commitments.

(Biconditional-2): In Charity, Jill is required to donate the money if and only if after making the donation, she will still have a happy life.

(Biconditional-3): In Promise, Jack is required to pay his friend if and only if the date is bald.

To fully see the problem, let us be a little more precise with our terminology. We start by disentangling two distinctions. First, there is the distinction between *vague* and *precise* terms: ‘tall’ and ‘bald’ are vague, while ‘prime’ is precise (and one can extend this notion to sentences, where any sentence containing a vague term is also vague). Second, there is distinction between *borderline* and *definite* statements.^{26, 27} Any borderline sentence must contain at least one vague term in it, but the converse isn’t the case. If Harry is two meters tall, then ‘Harry is tall’ is definite rather than borderline, and even when Harry is a borderline case of tallness, ‘Harry is tall or Harry is not tall’ is definite.

Now, consider a point in the development of the entity in the incubator where it is borderline whether it is a person.²⁸ Or consider a certain amount of money where, if Jill donates that amount, it is borderline whether she will still have a happy life; Or consider a case where the date has just a little bit of hair so it is borderline whether they are bald.

Since the right-hand statement in each of these biconditionals is borderline, it might seem to trivially entail that the left-hand statement is borderline as well. Moreover, since it seems that the relevant source of borderliness in each of these cases is the moral term, these moral terms must be vague.²⁹

This line of argument is too quick: no epistemicist should accept that if ‘q’ is borderline and ‘q if and only if p’ is true, then ‘p’ must be borderline as well. After all, consider a case where ‘Harry is tall’ is borderline. The epistemicist will nevertheless accept that there is some true biconditional of the form ‘Harry is tall if and only if Harry is over *n* meters in height’, i.e. a biconditional where the left-hand side is borderline, but the right-hand side is phrased using of fully precise language, and hence is not borderline. However, note that such cases are ones where the biconditional in question is one where we do not and cannot know its truth, and very plausibly the biconditional should itself be classified as borderline.

A much more compelling principle, one that most epistemicist *do* accept, is the following.³⁰

Distribution: if ‘q’ is borderline and ‘q if and only if p’ is definite, then ‘p’ is borderline.

²⁶ As standard, ‘p’ is definite iff it true and is not borderline, and conversely, ‘p’ is borderline if and only if neither it nor its negation is definite. Note that this means that some statements are neither definite nor borderline.

²⁷ There is a delicate question of what kind of entity ‘borderline’ and ‘definite’ apply to. Strictly speaking, I take them to apply to utterances or sentences in context (but informally and where the context is clear I will sometimes talk of them as applying to statements, propositions, or sentences).

²⁸ Of course, there is a view according to which the entity is definitely a person from the point of conception. But those who take this view would either reject Biconditional-1 or reject the claim that Incubator is a case of moral vagueness. In that case, as mentioned above, a different example can be used throughout the discussion.

²⁹ Obviously, other terms in the moral statement might be vague as well (‘pay’, ‘destroy’ and so on). However, we can easily set up the details of each case so that none of the other terms are at risk of rendering the whole moral statement borderline (e.g. assume that the way the entity in the incubator is destroyed is a very clear case of ‘destruction’ and so on).

³⁰ The principle is called ‘Distribution’ because it is equivalent to the claim that ‘definitely’ distributes over a biconditional. For explicit endorsement of Distribution in the epistemicist context see Williamson (1997, 271), Hawthorne (2006, 202), and Yli-Vakkuri (2016, 830). The only explicit rejection of it I’m aware of is in Magidor (2018).

Moreover, each of the three biconditionals are such that if they are true, they are arguably also *knowable*.³¹ And on nearly every theory of vagueness, one cannot know borderline statements (particularly so on the epistemicist theory, which takes borderlineness to be a form of ignorance). So, if the biconditionals are knowable, they are themselves definite.

We can summarise the problem as follows: there are a range of cases where we have a biconditional of the form '*m* if and only if *d*', where '*m*' is a moral statement and '*d*' a descriptive borderline statement, and where we have strong reasons to think '*m* if and only if *d*' is knowable. That is, we wish to accept the following (using 'B' and 'D' to denote borderline and definitely):

- (i) B('d')
- (ii) It is knowable that *m* if and only if *d*.

Since no one can know a borderline statement, this entails:

- (iii) D('m if and only if d')

Now, given Distribution (iii) entails the following:

- (iv) B('m')

But if '*m*' is borderline, it must contain at least one vague term. And since the only relevantly vague terms in the statements are the moral terms, these moral terms must themselves be vague. The upshot is that the vagueness of the descriptive statement seems to inevitably infect the moral statements, and renders the rejection of Moral Vagueness untenable for the epistemicist.³²

§4. How to stop the spread of the infection

How can the epistemicist respond to the problem of infectious vagueness? If we hold fixed the rejection of Moral Vagueness and hence maintain that the relevant moral statements are definite, each instance of the problem can be addressed by rejecting one of the following three claims:

- a. D('d iff m')
- b. B('d')
- c. B('d') ∧ D('d iff m') → B('m'), i.e. the relevant instances of Distribution.

In fact, I think different examples merit different reactions. For example, some cases we might be less tempted to think that the biconditionals are knowable, and hence have little reason to insist that they are definite, making the first option attractive. I will briefly discuss some reasons to opt for the second option in §4.4 below. But my main focus here will be on the third option: namely, rejecting the relevant instances of Distribution.

In Magidor (2018), I argued that – despite being commonly accepted – epistemicist should reject Distribution: some definite biconditionals have one side which is definite and one which is borderline. Of course, if the biconditional is itself known, then if the definite side (the moral statement, in our

³¹ Should the fact that these principles are controversial mean that they are not knowable even if true? That strikes me as veering towards quite widespread scepticism about moral knowledge. If one is happy to accept, (as I think we should), that moral knowledge is reasonably common then I don't see why we couldn't have moral knowledge of principles such as the three biconditionals.

³² It is worth noting that the problem of infectious vagueness isn't in any way particular to epistemicism: most views of vagueness accept Distribution (See Magidor 2018, 145). However, for some views of vagueness, it might be less troubling to accept that moral terms are vague (or at least, it will not be troubling for the reason outlined in §1). And at any rate, my response to the problem of infectious vagueness is specific to Epistemicism, so I will restrict the discussion to this view.

case) were to be known, that would entail that the other side (the descriptive statement, in our case) would be knowable as well, which contradicts the claim that it is borderline. This means that the definite statement must be *unknowable*. But crucially, this does not entail that the moral statement is borderline: as I stressed in §2, being borderline isn't the only possible form of ignorance: a claim might be unknowable for other reasons. In our case, this observation allows for the possibility that the moral statement is not vague, and hence that moral terms are not semantically plastic.

While in Magidor (2018) I argued that Distribution in its full generality fails, what we need here isn't merely the claim that Distribution has *some* counterexamples, but rather that it can have the counterexamples of the specific form that is relevant to this case.

My plan is to offer some models for how it might come about that biconditionals such as 1-3 can be definite despite involving a descriptive claim which is borderline and a moral claim which is definite and indeed not even semantic plastic. To be clear, I do not argue that these models represent a realistic (meta)semantics for how the specific terms in question actually works in English. Rather, as the label 'models' suggests, they serve to illustrate various ways in which the vagueness of the descriptive claims doesn't have to infect the relevant moral claims, and thus that epistemicism without moral vagueness is a coherent view that at least has some in-principle resources to address the problem of infectious vagueness. In §4.1, I set the stage by spelling out in more detail what, according to the current framework, it takes for a statement to be borderline. In §4.2, I explain in fairly general terms what sort of structures the models we are looking for might exhibit. In §4.3, I provide some more concrete examples for how the structures in question could be realised. Finally, in §4.4 I briefly discuss other routes for blocking the problem of infectious vagueness.

§4.1 Condition for borderliness

Before we proceed, we need to be a bit more precise about what, on the current framework, it takes for an utterance to be *borderline*.

According to Williamson (1994), even when we have a true borderline statement, we do not and cannot know that it is true. Moreover, there is a specific type of explanation for this ignorance of borderline statements in terms of semantic plasticity.

Suppose for example that the actual sharp cut-off for 'tall' is 1.8m, and that Jill's height is 1.805. Suppose her height is close enough to the cut-off for 'tall' to render 'Jill is tall' borderline (of course, since Jill's height is over the cut-off for 'tall', it is a *true* borderline statement). Now consider an agent in the actual world, who believes that Jill is tall. We want to explain why this true belief fails to constitute knowledge.

Williamson's explanation goes roughly like this: since 'tall' is semantically plastic, there is a close world *w* in which the use-facts for 'tall' are slightly different and consequently 'tall' has a slightly different meaning: call it 'tall*', which – let us suppose - has the cut-off height of 1.81m. Now, since speakers are presumably not sensitive to these small differences in meaning, in *w*, the agent will have a corresponding belief to their actual one, a belief that they express in exactly the same way as in the actual world, i.e. by endorsing 'Jill is tall'. However, due to the semantic plasticity of 'tall', this *w*-belief would have a slightly different content than its actual content (namely, that Jill is tall*). And moreover, that content is false in *w*: we're holding Jill's height fixed across close worlds, so in *w* Jill is still 1.805m in height, which is under the cut-off for being tall*. This means that in *w*, the agent has a false belief

that is very similar to their actual true belief. This fact, according to Williamson, is sufficient to render the agent's actual belief as epistemically *unsafe* and thus it doesn't constitute knowledge.³³

Putting this a bit more formally, we might propose the following sufficient condition for being borderline.³⁴

Condition for Borderlineness:

An actually true sentence 'p' is borderline if for every speaker S, there is a close possible world w where the following three conditions hold:

Plasticity: due to its semantic plasticity, 'p' has a slightly different meaning in w than it actually has.

Falsity: The content expressed by 'p' in w is false in w.

Doxastic Similarity: If in the actual world S forms a (true) belief that they express by 'p' then in w they have a corresponding (false) belief that they also express by 'p'.

With this account in place, we can now consider what sort of models the epistemicist can provide that block the problem of infectious vagueness.

§4.2 Two structures for potential models

Recall that our aim is to show that there are models where the following structure is realised:

- (i) Borderlineness of the descriptive claim: '*d*' is borderline.
- (ii) Definiteness of the moral claim: '*m*' is definite.
Moreover, recalling that our motivation for thinking that the moral statement is definite was to avoid the semantic plasticity of moral terms, we will look for models where the relevant moral terms are not even semantically plastic.
- (iii) Definiteness of the biconditional: '*m* if and only *d*' is definite.

³³ Note that this relies on an extended notion of safety, on which one can violate safety not only when one's very belief could have easily been a false belief but also when one could have a sufficiently similar (though distinct) false belief. As Williamson points out, some sort of extended notion of safety is anyhow required to explain, for example, why believing true mathematical statement because of a lucky guess does not constitute knowledge (Williamson 2016, 840). For a critical discussion of the notion of safety in play see Kearns & Magidor (2008).

³⁴ See Magidor (2018) for a discussion of some of the subtleties in phrasing this condition. Most importantly, I argue there that the condition as phrased above is too weak. In a nutshell, the problem (also raised in Caie 2012) is that as phrased, nothing ensures that the underlying conditions are held fixed in w: suppose, for example, that Jill is actually 1.9m in height and that this is sufficient to render 'Jill is tall' definite. The problem is that there might be a close world w where the cut-off for 'tall' is 1.81, but also where Jill is slightly shorter than she actually is – for example where she is only 1.8m in height. As long as w is sufficiently close, this would satisfy the above condition and wrongly predict that 'Jill is tall' is borderline.

To address this, I propose strengthening the sufficient condition for borderlineness, by ensuring that the ultimate grounds of the truth of the statement in the actual world are held fixed.

As it turns out, though, this complication does not matter for our current purposes: in all models discussed in §4.3 the underlying conditions for the statement we wish to show is borderline will not vary across the worlds discussed (that is, facts like the organism's biological stage of development, Jill's financial situation, or the number of hairs on the date's head will all be held fixed). So for simplicity, I keep the condition in the simplified form given above.

Moreover, recalling that our main motivation for thinking that the biconditional is definite was to allow it to be knowable, we will look for models where some agent *S* knows that the biconditional is true.

What form can these models take?

First note that all models must be ones where, all three claims '*m*', '*d*' and '*m* if and only if *d*' are true in the actual world. (The moral claim and the biconditional must be true because they are definite, and the descriptive claim must be true because, by the true biconditional, it is materially equivalent to a true claim.)

In order to ensure the borderlineness of the descriptive claim, we utilise Condition for Borderlineness, and assume that there is a possible world *w* which satisfies all three clauses (*Plasticity*, *Falsity*, and *Doxastic Similarity*) with respect to '*d*' and some agent *S* – an agent that, on our models, knows the biconditional.

The tricky part will be to ensure that the existence of *w* does not threaten the truth of (ii) or (iii) above. There are two potential structures that enable this.

According to the first structure, '*m* iff *d*' expresses a true proposition not only in the actual world, but also in all close worlds (that is, for each close world *v*, the biconditional expresses in *v* a proposition that is true in *v*). This makes knowing the biconditional easy: *S* isn't at risk of having a false belief in the biconditional. (iii) is thus easily satisfied.

What about (ii) though? Recall that according to *Falsity*, '*d*' expresses in *w* a proposition that is false in *w*. Moreover, since we are assuming the biconditional expresses a truth in all close worlds, then '*m* if and only if *d*' also expresses a truth in *w*. But assuming that the connective 'if and only if' has a stable meaning across all close worlds (i.e. it expresses the ordinary truth-functional biconditional), this entails that in *w*, '*m*' must express in *w* a claim that is false in *w* as well.

Does the fact that '*m*' expresses a truth in the actual world and a falsehood in *w* entail that '*m*' is semantically plastic, thus violating (ii)? The answer is negative. Suppose that in the actual world, I happen to have precisely 100,000 hairs on my head, but in a nearby world *w*, I happen to have 100,001 hairs on my head. The sentence 'Ofra has precisely 100,000 hairs on her head' is true in the actual world, and false in *w*. But we can explain this without any appeal to semantic plasticity: the sentence means precisely the same in both worlds – it's just that the conditions relative to which the claim is evaluated are different.³⁵ There is thus scope for a view on which '*m*' is not semantically plastic (and ipso facto, not borderline), yet expresses a falsehood in *w* simply because the situation relative to which it is evaluated is relevantly different. (More on the specifics of how this might happen in the next section).

On our first structure, the moral claim '*m*' is false in the close world *w*. By contrast, the second structure is one on which '*m*' expresses a truth in all close worlds. This makes it easy to accommodate (ii), i.e. the claim that '*m*' is definite. But now we face a different challenge: in the close world *w*, '*d*' is false and, '*m*' is true (because it is true in all close worlds). This entails that the biconditional '*d* if and only if *m*' expresses a falsehood in *w*. Moreover, since '*d*' is semantically plastic, the biconditional is semantically plastic too.

³⁵ This important observation was first highlighted by Hawthorne (2006) who distinguishes between semantic plasticity and what he calls 'extensional plasticity' (namely, a sentence or predicate having a different extension in a close world simply because the relevant conditions have changed).

Does the fact that the biconditional is semantically plastic and expresses a falsehood in a close world entail that it is borderline, thus violating (iii)? The answer is negative. For Condition for Borderlineness to apply, we must also satisfy its third clause, namely *Doxastic Similarity*. But if there is some principled reason why our agent S does not believe the biconditional in close worlds in which it is false (including w), there is no barrier for S to know that biconditional or for the biconditional to be definite.

To summarise: there are in principle two types of models that can accommodate (i)-(iii). On both types of models there is a close world w where, due to semantic plasticity, the descriptive claim is false. On Type 1 models, the biconditional is true in w and the moral claim is false, but the moral claim nevertheless expresses the same proposition that it expresses in the actual world (and hence the moral claim doesn't satisfy *Plasticity*). On Type 2 models, the moral claim is true in w, and the biconditional is false, but not falsely believed by the relevant agent S in any close worlds (hence the biconditional doesn't satisfy *Doxastic Similarity*).³⁶ In the next section I sketch out some toy models that can realise these two types.

§4.3 Some specific models

In this section I outline three toy models that realise these structures. To reiterate the above caveat: these do not purport to be a realistic account of the way these terms actually work in English. Rather, they are intended as a demonstration of feasibility, illustrating various ways our language could in principle work so as to allow descriptive vagueness without the corresponding moral vagueness.

Model 1 (Type 1, Promise): Suppose that in the actual world, Jack makes a promise to his friend using the following utterance 'I promise to give you ten pounds if your date is bald'.

Assume that in the actual world, the date has exactly 10 hairs and 'bald' picks out the property of *having 10 hairs or less*. Thus 'the date is bald' is true in the actual world. Suppose that due to the semantic plasticity of 'bald' there is a close world w in which 'bald' picks out a slightly different property (call it 'bald*'), which is the property of *having 9 hairs or less*. Thus in w, 'The date is bald' expresses the claim that the date is bald*, i.e. that they have 9 hairs or less. But since in w, the date also has exactly 10 hairs³⁷, 'The date is bald' expresses in w a claim that is false in w.

Now, assume that, Jack utters the promise sentence ('I promise to give you ten pounds if your date is bald') not only in the actual world, but in all close worlds. Since people are required to keep their promises in all worlds (certainly in all close ones), the biconditional 'Jack is required to give ten pounds if and only if the date is bald' expresses a true claim in all close worlds, including w. However, in w, the right-hand side of the biconditional expresses a falsehood, so the left-hand side is false as well: that is, 'Jack is required to give his friend ten pounds' is false in w.

However, 'Jack is required to give his friend ten pounds' is not semantically plastic despite being true in the actual world and false in w. It expresses precisely the same moral statement in both the actual world and in w. The difference in truth-value is because the underlying facts are crucially different

³⁶ We can now also see in more detail how the discussion can be reframed along the lines proposed in n.4. The *problem of infectious plasticity* can be summarised by the following argument: (i) the descriptive claim is a paradigmatic case of borderlineness: it is semantically plastic in a way that renders it false in a nearby world w; (ii) The biconditional is known, and hence true in all close worlds; So (iii) The moral statement must be false in the nearby world w. So, (iv) The moral statement must be plastic as well.

Type 1 models show how we can resist the argument by rejecting (iv), i.e. they show the moral statement can be false in close worlds despite not being plastic. Type 2 models show how we can resist the argument by rejecting (ii), i.e. they show how the biconditional can be known despite being false in nearby worlds.

³⁷ I assume the number of hairs is held fixed across close worlds, cf. n.34 above.

across the two worlds: in the actual world, Jack promised to give money if the date is bald; but in *w*, Jack made a promise with a slightly different content, namely to give money if the date is bald*. Moreover, since the condition of his *w*-promise aren't fulfilled in *w* (the date is not bald*), Jack is not required to pay in *w*, and the moral claim is no longer true in *w*.

Model 2 (Type-1, Charity): Model 2 realises the same general structure (namely one where the biconditional expresses a truth in all close worlds including *w*, and the moral claim is false in *w* despite expressing the same proposition as in the actual world). But this time the structure will be realised for slightly more elaborate reasons.

Let us start by telling a story about the metasemantics of 'happy'. There are multiple dimensions that feed into what we count as a 'happy life'. For simplicity let us assume there are just two such dimensions: pleasure and meaningfulness. These dimensions can pull in different directions (sometimes to have a more meaningful life, one must invest a lot of unpleasurable time and effort), and let us suppose the vagueness of 'happy' is due to the fact that it is not obvious how exactly these factors are weighted: 'happy' in the actual world picks out a property that requires a certain weighting of these factors, while in other close worlds, the term picks out a property with a slightly different weighting.

Moreover, let us suppose that how these dimensions are weighted is due to a certain stable *functional role* that the term 'happy' plays: in all close worlds, it is used by the relevant linguistic community to describe the kind of life that has the positive-making features that are valued in that community. In a community which values pleasure over meaningfulness, that functional role is best fulfilled by a property that gives higher weighting to pleasure, and vice versa in a community that values meaningfulness over pleasure. In particular, assume that in the actual world, the community places a higher value on meaningfulness than in *w*, and correspondingly in the actual world 'happy' picks out a property that gives a higher weighting to meaningfulness, while in *w* it picks out a slightly different property (*happy**) that give a higher weighting to pleasure.

Now assume that after Jill donates the money required to save her colleague's life, her own life is slightly less pleasurable but slightly more meaningful. Assume that after donating her life is thus still happy, but no longer happy*. Thus, due to the semantic plasticity of 'happy', the descriptive claim 'after donating, Jill's life is still happy' expresses a truth in the actual world and a falsehood in *w*.

Finally, let us add some toy-moral theory to our model. Suppose that whether Jill is required to donate, depends at least in part on the value people in her community place on her post-donation situation: she isn't required to donate at all costs, and how we evaluate the relevant "costs" depends – at least in part – on subjective measures of her post-donation situation. Crucially, note that this toy moral theory is not in tension with moderate moral realism (cf. note 2 above). It is not that *in general* what we are required to do depends on what members of the community subjectively value. It's rather that in this particular case, the moral requirement is essentially balancing the needs of her colleague and her own wishes and needs – so for these purposes it matters what these needs are. (This is thus no different from the fact that the question of whether you are morally required to share your ice cream with a friend depends at least in part on the question of whether your friend likes ice cream.)³⁸

Assume that as an upshot, in the actual world where the community places a higher-value on meaningfulness, the effective "cost" to Jill is lower, and she is required to donate. By contrast, in *w*, where members of the community place a higher value on pleasure, the "cost" to Jill is higher, and

³⁸ Cf. Schoenfield (2016, n. 13), who also notes that moral realism is consistent with some moral facts depend on social facts.

she is not required to donate. The moral claim is thus true in the actual world and false in w : not because 'required' is semantically plastic, but rather because the relevant moral conditions differ. And given the different interpretations of 'happy' in each world, the biconditional 'Jill is required to donate iff she still has a happy life' is still true in both worlds.

The idea underlying this second model is ultimately this: the fact that the community in w weighs pleasure and meaningfulness slightly differently than in the actual world, is a common factor in determining two very different kinds of facts: first, it plays an indirect role in the metasemantics of 'happy', ensuring the term picks out a different property in each world. And second, it affects the moral requirements – which (in this particular case) depend at least to some extent on what members of the relevant community value. This common factor allows the extension of 'required' and 'happy' to covary in a way that ensures that the relevant biconditional is true across close worlds.

Model 3 (Type 2, Incubator):

The third and final model is one which realises the second structure discussed in §4.2, namely one where the moral claim expresses a truth in all close worlds, the biconditional is false in some close worlds, but where the relevant speaker S does not believe it in such worlds.

Let's start with the moral facts: in this model, we suppose that whether or not it is permissible to destroy the contents of the incubator supervenes entirely on the entity's biological stage of development. And suppose that in the actual world, the entity in the incubator is at a sufficiently early stage of development that it is (just about) permissible to destroy it. Moreover, since we're holding fixed the entity's stage of development across all close worlds, this moral fact also obtains across all close worlds. And finally, since 'It is permissible to destroy' is not semantically plastic, the moral sentence expresses a truth in all close worlds.

Next, let us say something about the metasemantics of 'person' in the actual world. The meaning of the term depends on a range of use facts, but let us suppose that one use fact that played a particularly crucial role is that many years ago, a famous supreme court ruling stated: 'It is permissible to destroy an entity in an incubator if and only if it is not a person'. Indeed, this court ruling played such a central role in the metasemantics of 'person', that it is a constraint on the (current) meaning of the term in the actual world that that it must render this biconditional true. Since, as we have assumed above, it is in fact permissible to destroy the entity in our particular incubator, then 'person' picks out a property which the relevant entity does not satisfy and correspondingly, 'the entity is not a person' expresses a truth in the actual world.

Of course, if the same court ruling happens (and plays the same key role) across all close worlds, 'person' would be fixed in a similar way in all such worlds and would not be semantically plastic (at least not in a way that's relevant to our specific descriptive claim). We will return to this possibility in §4.4 below.

But for the purpose of our current model assume that in some close worlds, this court ruling did not happen. In those worlds, the metasemantics of 'person' need not be tied to the permission to destroy. In particular, assume that in some such close world w , 'person' picks out the property person^* , which the entity in our incubator *does* satisfy (i.e. in w , 'the entity is not a person' expresses a falsehood). This also means that in w , the biconditional 'It is permissible to destroy the entity if and only if it is not a person' expresses a falsehood.

Finally, assume that our agent S believes the biconditional in the actual world, but their only reason for believing it is that they have heard of the court ruling and take supreme court judges to be excellent

moral authorities. Correspondingly, S only believes the biconditionals in close worlds in which the ruling took place (and is considered similarly important). But recall that any such world is one in which the ruling plays the relevant metasemantic role, ensuring that the biconditional expresses a truth. Contraposing, any world in which the biconditional expresses a falsehood (including the world *w* just considered), is one where the ruling does not take place, and hence where S does not believe the biconditional. In other words, S is at no risk of having a false belief expressed using the biconditional, and hence there is no barrier to the biconditional being known by S and hence definite.

Note that while in Model 2 we had a common factor (what community members value) that affected both the metasemantics of the descriptive claim and the moral facts, this time we have a different kind of common factor. In this model, the supreme court ruling serves as a common factor in determining on the one hand, the metasemantics of 'person', and on the other hand, the doxastic attitudes of our speaker S. This common factor allows the two to co-vary in a way that ensures that the biconditional is not falsely believed across close worlds.³⁹

§4.4 other ways to block the infection

Even if one is convinced by my argument that at least in many cases, it is possible for there to be failures of Distribution of the relevant kind, one might worry that these structures will not necessarily be realised in every purported case of infectious vagueness. Since a single case of infectious vagueness would be sufficient to render moral terms vague, that would indeed be a problem for a view that they are not.

I agree that it is highly implausible that every alleged case of infectious vagueness would make for a failure of Distribution. However, recall that for each potential biconditional of the form '*m* if and only if *d*' there are two additional options for addressing the problem of infectious vagueness. The first, to maintain that the descriptive claim '*d*' is not, after all, borderline; the second, to maintain that the biconditional is not, after all, definite.

If we hold as a constraint that moral properties are reference magnets and the moral terms are thus not semantically plastic, these two options might not be as unattractive as they might initially seem.

Return to the case of Incubator. Suppose that, as suggested in the discussion of Model 3, in all close worlds the moral claim is true. But suppose that this time the metasemantic constraint that the biconditional ('permissible to destroy iff is a person') must express a truth not only in the actual world, but also in all nearby worlds. (This could happen, for example, if the relevant court ruling in Model 3 happens in all close worlds.). Since the moral claim is also true in all close worlds, the descriptive claim 'the person in the incubator is a person' will be true in all such close worlds, and the sentence will not be borderline. Indeed, on this view, 'person' will not even be semantically plastic (at least not along the relevant dimension). But this hardly surprising: if moral terms have a stable meaning, and the meaning of 'person' is directly tied to that of a moral term, then the latter would be stable as well. In other words, instead of a case of infectious vagueness (where the vagueness of the descriptive term

³⁹ There is a more general lesson here for the connection between semantic plasticity and ignorance. Williamson suggests that semantic plasticity leads to ignorance because speakers' beliefs aren't sensitive enough to track the truth of a plastic claim and thus ensure they avoid false beliefs in nearby environments. But in our current case, speakers' beliefs do succeed in tracking the truth of the semantically plastic biconditional, not because they have some special insight into what propositions it expresses and when that proposition is true but rather because of a common factor that plays a role both in what they believe and in what proposition that belief happens to express. Thus, there can be subtle ways in which plasticity might sometimes resist ignorance.

leads the moral term to be vague), we have a case of *infectious precision* (where the precision of the moral term leads to the precision of the descriptive one).

Alternatively, suppose that the biconditional is false in some close worlds, but that unlike in Model 3, there is no systematic reason that speakers don't believe it all such worlds. For each speaker, there are some close worlds where they falsely believe the biconditional. Of course, if one falsely believes the biconditional in a close world, then any actual (true) belief in it isn't safe, and hence cannot constitute knowledge. But then there is no longer pressure to take the biconditional to be definite (recall that the argument for why the biconditional must be definite relied on the claim that it is knowable). Put in less technical terms: if we stick to our belief in the biconditional, ignoring the fact that the moral statements are stable (in both their meanings and the underlying moral reality), while the descriptive statements are shifty, then it is hardly surprising that our belief in the biconditional cannot constitute knowledge.

And finally, remember that there is no pressure to resolve all cases of alleged infectious vagueness in the same manner: it is perfectly plausible that in some cases the biconditional is not after all known; in others the descriptive claim is not after all borderline; while in others we get failures of Distribution. Either way, the view that maintains Epistemicism and Moral Realism, yet rejects Moral Vagueness remains an attractive and viable way to escape the trilemma posed by Incompatibility.⁴⁰

⁴⁰ Thanks to S. Matthew Liao, David Liebesman, and Timothy Williamson as well as to audiences at, PhLip, 'Moral Vagueness' workshop in Upsala, and the Lauener Symposium in Bern for helpful discussions.

References

- Abasnezhad, A. (2023), 'Moral realism and semantic accounts of moral vagueness', *Inquiry* 66: 381-393.
- Balcerak Jackson, B. (2020), 'Verbal disputes in metaphysics', in Ricki Bliss & J.T.M. Miller (eds.), *The Routledge Handbook of Metametaphysics*, London: Routledge, pp. 118-129.
- Caie, M. (2012), 'Vagueness and semantic indiscernibility', *Philosophical Studies* 160: 365-377.
- Constantinescu, C. (2014), 'Moral vagueness: A dilemma for non-naturalism', *Oxford Studies in Metaethics* 9: 152-185.
- Dougherty, T. (2014), 'Vague value', *Philosophy and Phenomenological Research* 87: 352-372.
- Dougherty, T. (2017), 'Vagueness and indeterminacy in ethics', in Tristram McPherson & David Plunkett (eds.), *The Routledge Handbook of Metaethics*, London: Routledge, pp. 185-193.
- Dunaway, B. & McPherson, T. (2016), 'Reference Magnetism as a Solution to the Moral Twin Earth Problem', *Ergo* 3: 639-675.
- Eklund, M. (2007), 'Characterizing Vagueness', *Philosophy Compass* 2: 896–909.
- Eklund, M. (2017), *Choosing Normative Concepts*, Oxford University Press.
- Hawthorne, J. (2006), 'Epistemicism and semantic plasticity', chapter 9 of his *Metaphysical Essays*, Oxford University Press.
- Hawthorne, J. (2022), 'Moral vagueness and epistemicism', *Journal of Ethics and Social Philosophy* 23: 214-247.
- Horgan, T. and Timmons, M. (1991), 'New wave moral realism meets moral twin earth', *Journal of Philosophical Research* 16:447-465.
- Kearns, S. and Magidor, O. (2008), 'Epistemicism about vagueness and metalinguistic safety', *Philosophical Perspectives* 22 (1):277-304.
- Lewis, D. (1983), 'New work for a theory of universals', *Australasian Journal of Philosophy* 61: 343-377.
- Liao, S. M. (ms.), 'Threshold deontology: lessons from vagueness' (unpublished manuscript)
- Magidor, O. (2018), 'Epistemicism, distribution, and the argument from vagueness', *Nous* 52: 144-170.
- Manley, D. (ms.), 'Moral Realism and Semantic Plasticity', (unpublished manuscript)
- Schoenfield, M. (2016), 'Moral vagueness is ontic vagueness', *Ethics* 126: 257-282.
- Wasserman, R. (2012), 'Personal identity, indeterminacy, and obligation' In Georg Gasser and Matthias Stefan (eds.), *Personal Identity: Complex or Simple?*, Cambridge: Cambridge University Press, pp. 63-81.
- Williams, J.R.G. (2018) 'Normative reference magnets', *Philosophical Review* 127: 41-71.
- Williamson, T. (1994), *Vagueness*, London: Routledge.

Williamson, T. (1997), 'Reply to commentators', *Philosophy and Phenomenological Research* 57: 945–953.

Williamson, T. (2000), *Knowledge and Its Limits*, Oxford: Oxford University Press.

Williamson, T. (2016), 'Reply to Yli-Vaakuri', *Canadian Journal of Philosophy* 46: 839-851.

Yli-Vakkuri, J. (2016), 'Epistemicism and Modality', *Canadian Journal of Philosophy* 46: 803-835.