

Vagueness and semantic indiscriminability

Michael Caie

© Springer Science+Business Media B.V. 2011

Abstract I argue, pace Timothy Williamson, that one cannot provide an adequate account of what it is for a case to be borderline by appealing to facts about our inability to discriminate our actual situation from nearby counterfactual situations in which our language use differs in subtle ways. I consider the two most natural ways of using such resources to provide an account of what it is for a case to be borderline and argue that both face crippling defects. I argue that the problems faced by these two accounts point to more general reasons to be skeptical of the claim that facts about semantic indiscriminability provide sufficient resources for an analysis of what it is for a case to be borderline.

Keywords Vagueness · Epistemicism · Metalinguistic safety · Timothy Williamson

1 Introduction

The phenomenon of vagueness is characterized by the existence of borderline cases. If it is, say, borderline whether Ben is bald then one cannot come to know that Ben is bald. In general, it being borderline whether ϕ provides a barrier to knowledge that ϕ . What is the nature of vagueness and of the borderline status in particular that precludes knowledge?

M. Caie (✉)
Department of Philosophy, University of California, Berkeley,
314 Moses Hall, Berkeley, CA 94720-2390, USA
e-mail: caie@berkeley.edu

One answer—that provided by standard semantic theories of vagueness—is that borderline cases are cases of truth-value gaps.¹ If it is borderline whether Ben is bald, then it is not true that Ben is bald nor is it true that Ben is not bald. Since knowledge that ϕ requires that it be true that ϕ , one cannot know that Ben is bald nor can one know that Ben is not bald.

An epistemic theory of vagueness holds that vagueness is compatible with bivalence. According to an epistemicist it may be both true that ϕ and yet also borderline whether ϕ . In borderline cases, then, there is something about which we are ignorant. Since an epistemicist does not think that borderline cases give rise to truth-value gaps, she must provide an alternative account of what it is for a case to be borderline that can explain our lack of knowledge in such cases.

In Williamson (1994), Timothy Williamson develops an account of vagueness that is intended to answer this demand. According to Williamson, “[t]he vagueness of an expression consists in the semantic differences between it and other possible expressions that would be indiscriminable by those who understood them” (Williamson 1994, p. 237). Although Williamson in that work and in other places has much to say about our inability to discriminate the actual semantic values that our words have from other possible semantic values that they might have had, and the relationship between our inability to make such discriminations and vagueness, he has never to my knowledge provided an explicit analysis of what it is for it to be borderline whether ϕ using these resources. Despite this it does seem that he thinks that these resources should suffice for this purpose. Thus in his response to a review by Vann McGee and Brian McLaughlin, Williamson complains that the theory favored by McGee and McLaughlin does not provide resources to explain what it is for a case to be definite. In contrast, he claims that “...the epistemicist can explain what it is for a to be definitely F as the absence of obstacles of a special kind to knowing that a is F...” (Williamson 2004 p. 118). ‘It is borderline whether...’ can be defined in terms of ‘It is definite that...’ as follows: It is borderline whether $\phi =_{df}$ It is neither definite that ϕ nor definite that $\neg\phi$. According to Williamson, then, the epistemicist can explain what it is for it to be borderline whether ϕ as the presence of a special obstacle to knowing whether ϕ or $\neg\phi$. The nature of this obstacle is in turn to be explained in terms of competent speakers’ insensitivity to slight differences between the semantic values that the terms of our language actually have and the semantic values that they would have given slightly altered global patterns of language use.

Despite Williamson’s assurances, I am skeptical whether the facts to which he appeals provide adequate resources for an account of what it is for a case to be borderline. The purpose of this paper is to develop the case for such skepticism.

2 Semantic indiscriminability and ignorance

According to Williamson, words and complex syntactic units such as sentences have classical semantic values. The semantic values of the words of our language are

¹ Truth-value gaps are postulated by standard supervaluationist and standard degree theories of vagueness. There are, however, other versions of semantic theories of vagueness, e.g., that proposed by Vann McGee and Brian McLaughlin, that do not clearly fit this mould.

fixed (at least in part) by the dispositions of speakers of the language to use sentences involving such words to make assertions, commands etc. The semantic values of most words are, according to Williamson, quite sensitive to slight changes in the overall pattern of use. The dispositions of individual speakers, however, to apply terms is less sensitive to such slight changes. Individual speakers' dispositions to use the terms of a language are, then, not sufficient to discriminate what the actual semantic values of the words of our language are from all the other possible semantic values that these words might have had. Let us label this phenomenon *semantic indiscriminability*.

Semantic indiscriminability provides, according to Williamson, a barrier to knowledge. Consider the predicate 'bald'. Which property this picks out is presumably quite sensitive to subtle facts about our dispositions to apply the term. Were we to use it slightly differently, certain individuals who satisfy 'bald' as we actually use it would not satisfy this predicate given our counterfactual use. Let 'Ben' name one of these individuals. While the sentence 'Ben is bald' expresses a truth, there is a close possible world in which it expresses a distinct proposition that is false at that world (assuming that the referent of 'Ben' does not change in this world). While the semantic value of a term such as 'bald' is sensitive to minute changes in the global pattern of language use, the individual dispositions of speakers to apply this term are not. If one is disposed to judge that Ben is bald given how things actually are, then one will be similarly disposed in a close possible world in which the facts about Ben's hair are held fixed but in which the patterns of applying the term 'bald' differ slightly. The disposition to judge that Ben is bald, then, will not in fact lead to a false judgment, but it easily could have. Such a judgment then cannot be an expression of knowledge.²

Let us grant for the purpose of this paper that Williamson is right that semantic indiscriminability provides a barrier to knowledge; indeed this seems to me quite plausible.³ As the above case illustrates, there are at least some cases in which we cannot come to know whether or not it is the case that ϕ due to semantic indiscriminability that are also cases in which intuitively it is borderline whether ϕ . And indeed the above case would seem to be an instance of a general pattern. For each vague predicate P it is plausible that there will be semantic indiscriminability of the kind exhibited by 'bald', and so there will be individuals such as Ben who are in some but not all of the indiscriminable extensions that the predicate might have had. In such a case it will intuitively be borderline whether the individual is P, and we will have an explanation for why in that particular case one cannot come to know that the individual is P. We have reason to think, then, that it might be possible to provide a general account of what it is for a case to be borderline by appeal to facts concerning our inability to discriminate actual from counterfactual semantic values, and that such an account would be able to explain why it is that we cannot come to know that ϕ when it is borderline whether ϕ .

As attractive as this thought is, I can't, however, see a way to make it work.

² For the full development of this idea see chapter 8 in Williamson (1994).

³ But see Kearnes and Magidor (2008) for some worries about whether metalinguistic safety is in fact necessary for a judgment to count as knowledge.

3 Some preliminary remarks

How can facts about our inability to discriminate actual from counterfactual semantic values be used to provide an account of what it is for a case to be borderline?

There is an obvious parallel between the facts in terms of which we are trying to provide an account of borderline cases and the facts in terms of which a supervaluationist analyzes what it is to be a borderline case. According to a supervaluationist, there are different ways in which our language could be made precise, i.e., there is a set of classical semantic assignments which count as admissible precisifications of our language. It is in terms of this set of classical semantic assignments that a supervaluationist provides an account of what it is for a case to be borderline. According to our epistemicist, there are a number of classical semantic assignments which are indiscriminable (in the sense indicated above) from the actual classical semantic assignment that our language has. It is in terms of this set of assignments that we will want to provide an account of what it is for a case to be borderline.

A supervaluationist model M is a quadruple $\langle D_m, \Delta_m, R_m, I_m \rangle$. D_m is a set of individuals. Δ_m is a set of points relative to which sentences are assigned truth-values. R_m is an accessibility relation that holds between members of Δ_m . I_m is an interpretation function which assigns to expressions classical extensions relative to members of Δ_m . Truth relative to a point is defined in the obvious way for atomic sentences. Similarly for Boolean combinations of sentences. A sentence of the form \lceil It is borderline whether $\phi \rceil$ is true relative to a point δ just in case ϕ is true for some δ' such that $\delta R_m \delta'$ and false for some δ'' such that $\delta R_m \delta''$.⁴

Obviously at this level of abstraction there is nothing distinctive about the supervaluationist account; for all that has been said 'it is borderline whether...' could have the same sense as 'it is contingent whether...'. What is distinctive about supervaluationism is the story about what truth at a point represents. According to the supervaluationist, for the intended model M there will be a bijection between the points in the model and the set of admissible precisifications. Under the bijection each point $\delta \in \Delta_m$ will be paired with some semantic assignment σ_δ which is an admissible precisification of our language. The pairing is such that every sentence ϕ will be true at δ just in case ϕ is true (at the actual world) according to σ_δ . According to the supervaluationist, it is borderline whether ϕ just in case there is an admissible precisification of the language in which ϕ is true and an admissible precisification of the language in which ϕ is false.

The natural thought is that just as the supervaluationist can provide an account of what it is for a case to be borderline by telling us how truth at a point is determined by the set of precisifications, so too can the epistemicist provide an account of what it is for a case to be borderline by telling us how truth at a point is determined by the set of semantic assignments indiscriminable from the actual assignment. This

⁴ The *locus classicus* for supervaluationism is Fine (1975). Note, however, that the formalism developed by Fine contains a number of additional features beyond those mentioned above. See also Keefe (2000) and Williamson (1994) for more detailed discussion of the supervaluationist framework.

thought would seem to be endorsed by Williamson. Thus, in the appendix to Williamson (1994) develops a model-theoretic framework within which the borderline operator can be interpreted which is essentially the same as a supervaluationist model.⁵ According to Williamson: “At least in its simplest form, epistemicism can take over and reinterpret the formal apparatus of supervaluationism, with a Tarskian conception of truth. To say that a valuation V^* is admissible by a valuation V is now to say that V^* is indiscriminable from V in the sense indicated” (Williamson 2003, p. 710).

There are two natural ways that I can see of using the set of indiscriminable semantic assignments to determine truth at a point. These correspond to two potential analyses of a case being borderline in terms of facts concerning semantic indiscriminability.⁶ I’ll now argue that neither of these provides an adequate account of what it is for a case to be borderline. The problems that these two analyses face will point to more general reasons to be skeptical that there is an analysis of what it is to be borderline in terms of the set of indiscriminable semantic assignments.

4 First analysis

What prevents us from knowing that Ben is bald or that Ben is not bald, according to Williamson, is the fact that in some nearby counterfactual situation in which our language use differs in imperceptible ways the sentence ‘Ben is bald’ expresses a proposition that is false at that world and in some other such situation ‘Ben is bald’ expresses a proposition true at that world. We want it to be the case that when it is borderline whether ϕ one cannot know that ϕ or that $\neg\phi$. It is natural, then, to think that a point of evaluation δ will tell us what truth-value ϕ has at one of the nearby use worlds w given its semantic value at w .⁷

If this is right then the following analysis of what it is to be a borderline case suggests itself. Let $W_{@}$ be the set of nearby possible worlds in which our global pattern of language use differs (if at all) from our actual use of language in ways to which the linguistic dispositions of individual speakers are not sensitive. We say:

⁵ Two points. First, the formal framework employed by a supervaluationist will not be *exactly* the same as that employed by an epistemicist. Of particular note is the fact that truth for the supervaluationist will be identified with truth at all points in the model, whereas for the epistemicist there will be a designated point that will represent the correct interpretation of the language. Second, it’s worth noting that Williamson in the appendix actually develops two models. The second is more complex than the first. Instead of a single accessibility relation over a set of points, it employs a set of such relations. Such added complexities will not matter for the problems that I will be raising, and so I will ignore them.

⁶ It is somewhat unclear to me which of these Williamson had in mind as a way of interpreting the model-theory he suggests. Since I think neither of these work I won’t be too concerned to sort out which Williamson intended. For the record, the second analysis that I will develop seems to me the one most naturally suggested by the above quote, since it is formally closest to standard supervaluationism. But the evidence here is far from unequivocal.

⁷ N.B. This way of using a classical semantic assignment to determine truth at a point differs from that employed by the supervaluationist. The latter uses a classical assignment σ to determine the truth value of ϕ at a point by evaluating what truth value ϕ has under σ at the actual world. In the next section I’ll consider an epistemicist account that parallels the supervaluationist in this respect.

(B) \ulcorner It is borderline whether $\phi \urcorner$ is true (at the actual world) just in case there is a world w_1 in $W_{@}$ such that the proposition expressed by ϕ at w_1 is true at w_1 , and a world w_2 in $W_{@}$ such that the proposition expressed by ϕ at w_2 is false at w_2 .

The appealing thought is that we can understand the barrier to knowledge that borderline cases pose in terms of the barrier to knowledge posed by semantic indiscriminability. According to (B) the barriers are one and the same. For a case to be borderline just is for there to be a barrier to knowledge whether ϕ as a result of semantic indiscriminability.

This analysis while prima facie attractive faces two serious problems.

The first problem is that this analysis does not easily generalize to give an account of what is required for \ulcorner It is borderline whether $\phi \urcorner$ to be true as evaluated at an arbitrary possible world. But ultimately we want an analysis of ‘It is borderline whether...’ to tell us this. For, of course, the borderline operator naturally embeds within standard metaphysical modal operators. So I could say truly: “Tom is definitely not bald, but it could have been borderline whether Tom is bald, after all he could have had fewer hairs.”

It obviously will not do to say that \ulcorner It is borderline whether $\phi \urcorner$ is true at a world w just in case there is a world w_1 in $W_{@}$ such that the proposition expressed by ϕ at w_1 is true at w_1 , and a world w_2 in $W_{@}$ such that the proposition expressed by ϕ at w_2 is false at w_2 , for then there would be no difference between what is actually borderline and what could be borderline.

The natural suggestion at this point is that we generalize the condition as follows. Let W_w be the set of worlds which stand to w as the set $W_{@}$ stands to the actual world i.e., the set of worlds whose language use is indiscriminable from that at w , but otherwise just like w . We say:

(B_w) \ulcorner It is borderline whether $\phi \urcorner$ is true at a world w just in case there is a world w_1 in W_w such that the proposition expressed by ϕ at w_1 is true at w_1 , and a world w_2 in W_w such that the proposition expressed by ϕ at w_2 is false at w_2 .

There is a serious problem with (B_w). When we are evaluating whether or not \ulcorner It is borderline whether $\phi \urcorner$ is true at a world w what we are after presumably, by epistemicist lights, is whether given *our* use of language, together with the facts about that world, things are borderline. Whether, e.g., Ben is borderline bald at a possible world w should not depend upon whether the speakers at this world use ‘bald’ differently from how we do—it should simply depend on facts about the distribution of hair on Ben’s head. (B_w) fails to respect this intuition—the propositions expressed by ϕ in W_w could be radically different from those expressed by ϕ in $W_{@}$, giving us absurd results such as that it is possible that a man with a full head of hair is definitely bald.

To fix this problem one might hold that W_w should consist of worlds containing speakers as in the set $W_{@}$ but otherwise like w . The problem with this suggestion is that we will not always be able to hold on to the relevant facts in w while changing the world to allow for linguistic communities such as those in $W_{@}$. To take a somewhat trivial, but sufficiently instructive, example, the world could be such that

there are definitely no speakers at all. We cannot, however, get this verdict if we require W_w to contain speech communities.

It seems to me that if we go along with (B) it is completely mysterious what in general we can say about the conditions for a case to be borderline at an arbitrary possible world. For not every such world will have worlds related to it in the same way that the worlds in $W_{@}$ are related to the actual world. But, crucially, despite such a world not having any other worlds related to it in the appropriate way it still makes sense to say that given the way things are in such a world certain things might be borderline others definite. A world in which there are no speakers may nonetheless still be a world in which it is, say, borderline whether a particular rock is red. The fact, then, that this account makes it mysterious what the conditions are for a case to be borderline at such worlds despite there clearly being such conditions counts as a serious strike against the proposal.

A second (related) problem with this analysis is that it looks like it will predict that intuitively definite cases are borderline. Let 'E₁', 'E₂', 'E₃' be the names of three electrons. Consider the sentence: 'E₂ is closer to E₁ than E₃'. Intuitively this sentence need not be in any way vague. It could express the same proposition at every world in $W_{@}$. Nonetheless there is no guarantee that this proposition will be true at every such world, or false at every such world. Nearby worlds could differ slightly in the spatial relations of the relevant electrons, some making the proposition true, others making it false. By (B) it follows that 'It is borderline whether E₂ is closer to E₁ than E₃' is true. This seems like the wrong result.⁸

In response to this worry, one may claim that we need to add as an additional condition that there be at least two worlds in $W_{@}$ in which ϕ expresses different propositions. Perhaps some response along this line could be made to deal with this particular example. I won't, however, bother to consider how this would look in detail. First, this response doesn't help with the first worry. Second, while such an adjustment might help with this particular case it would fail to address the deeper underlying problem that the case points to.

The problem is the following. Whether a case is borderline *in the actual world* depends on how things are *in the actual world*.⁹ Whether it is borderline whether a particular stone is red depends on the color of the stone, whether it is borderline whether John is tall depends on John's height, and so on.¹⁰ But according to (B) it is not just how things are in the actual world but how things are in all the other worlds in $W_{@}$ that matters for whether a case is borderline. In most of these worlds things will be similar enough that this will not make a difference to borderline claims. Slight changes in the overall pattern of language use will not make a difference to John's height, or the color of various objects etc. But you cannot change the overall pattern of use without changing countless other things. By making the assessment of

⁸ Thanks here to Cian Dorr for helpful discussion.

⁹ Obviously this is true for claims of the form: It is borderline whether ϕ , only so long a ϕ does not contain modal vocabulary. I'll take this restriction as implicit here and in what follows.

¹⁰ And in general, whether a case is borderline in w depends on how things are in w . It is, I suggest, because it is how things are in a world w that matters for the assessment of borderline claims at w that we can assess for an arbitrary world whether it is borderline whether ϕ at w despite there not being worlds related to w as the worlds in $W_{@}$ are related to the actual world.

borderline claims depend on how things are at other possible worlds differing from the actual world in these respects one leaves oneself open to the possibility of getting the wrong result due to such differences. At best, then, (B) provides a useful heuristic. It cannot, however, provide an account of what it is for a case to be borderline.

5 Second analysis

In response to the above worries, it is natural to adjust (B) so that instead of evaluating each indiscriminable proposition at the world at which it is expressed by ϕ one evaluates all such propositions at the actual world. Let $W_{@}$ be as above. We say:

(B*) \ulcorner It is borderline whether ϕ^{\ulcorner} is true (at the actual world) just in case there is a world w_1 in $W_{@}$ such that the proposition expressed by ϕ at w_1 is true (at the actual world), and a world w_2 in $W_{@}$ such that the proposition expressed by ϕ at w_2 is false (at the actual world).

Note that this alternative avoids both of the objections leveled against (B).

If ‘ E_2 is closer to E_1 than E_3 ’ expresses the same proposition at all the worlds in $W_{@}$ then ‘It is borderline whether E_2 is closer to E_1 than E_3 ’ will come out false.

(B*) also has the virtue of smoothly generalizing to give truth-conditions for arbitrary possible worlds. We say:

(B*_w) \ulcorner It is borderline whether ϕ^{\ulcorner} is true at a world w just in case there is a world w_1 in $W_{@}$ such that the proposition expressed by ϕ at w_1 is true at w , and a world w_2 in $W_{@}$ such that the proposition expressed by ϕ at w_2 is false at w .

Despite these virtues, however, (B*) is inadequate. We were motivated to tie the truth of ϕ at a point to the truth of the proposition expressed by ϕ at a world *as evaluated at that world* because it is the truth value of such propositions at such worlds that is relevant to the safety of our judgments. We have seen, however, that this way of understanding truth at a point brings with it serious problems, and that these problems motivate an account according to which the set of indiscriminable propositions are all evaluated relative to the actual world. In making this change, however, the worry arises that we run the risk of severing the connection between a case being borderline and our judgments lacking safety. The worry is a good one. What I’ll now argue is that it is not in general true, given this analysis, that when it is borderline whether ϕ one does not know that ϕ .

The argument depends on two claims:

- (1) Given the above analysis, it is borderline whether ‘Everest’ refers to Everest.
- (2) It is known that ‘Everest’ refers to Everest.

No argument will be offered in support of (2). The claim that we know that ‘Everest’ refers to Everest is as solid as any claim that I could muster in its support. As a dialectical point, though, let me note that an epistemicist in particular is very unlikely to balk at this claim. For much of the support for the epistemicist theory

comes from an appeal to disquotational principles governing truth. But it would be very strange to hold that we know that ‘true’ behaves in a disquotational manner, but we do not know that ‘refers’ does as well.

It will help to have a label for the sentence “ ‘Everest’ refers to Everest”, so let’s call it ‘(E)’. To show that (I) holds we need to show (a) that there is a proposition expressed by (E) in a nearby possible world which is true at the actual world, and (b) that there is a proposition expressed by (E) in a nearby possible world which is false at the actual world.

(a) is easy. The actual world is one of the nearby possible worlds, and the proposition actually expressed by (E) is true at the actual world.

Turning to (b). The proposition expressed by (E) is a function of the semantic values of, ‘Everest’, ‘Everest’ and ‘refers’. I make the following assumptions about the semantic values of these terms as used at a possible world. I assume that the semantic value of ‘refers’ as used at a possible world will be an intension which maps a possible world to a relation, i.e., a set of ordered-pairs of objects. I assume that the semantic value of ‘Everest’ as used at a possible world will be an object.¹¹ Finally, I assume that there is no relevant vagueness in the quote name ‘Everest’. The term has as a semantic value the same word-type given our use of language in a nearby possible world w as it does given our actual use of language.¹²

Given this we can say the following. If the proposition expressed by (E) as used at some nearby possible world w is to be true at the actual world then it must be the case that:

(INT) The intension of ‘refers’ as used at w maps the actual world to a relation that holds between the word ‘Everest’ and whatever is the referent of the word ‘Everest’ as used at w .

I will now argue that not every nearby possible world will satisfy (INT). This suffices to establish the existence of a nearby possible world where our pattern of language use differs in indiscriminable ways and in which (E) expresses a proposition that is false at the actual world. This, then, suffices to establish (b). There are two arguments that I’ll give for this.

Argument 1 The first argument depends on the claim that there is a nearby possible world w_c where the following hold:

¹¹ Both of these assumptions could be altered in certain ways to accommodate alternative semantic frameworks. It wouldn’t matter if the semantic value of ‘refers’ were simply something which determined an intension of the kind in question, but was perhaps more fine-grained. Nor would it matter if the semantic value of ‘Everest’ were simply something which rigidly picked out an object, but was not identified with the object itself.

¹² Two points. First, note that the notion of word-type being employed here is one according to which a word-type may have different semantic values at different possible worlds. Assuming such a conception of word-type is fine in this context since this is the notion of word-type that is required to frame the epistemicist account. See Williamson (1994) chapter 8, footnote 12 for an explicit endorsement of this notion of word-type. Second, if you are worried that there may be some minor vagueness in this term, note that the arguments that follow could be reframed using the weaker and altogether plausible assumption that there are some nearby worlds in which the term ‘Everest’ refers to the same word-type as it does given our actual use of language, but in which that word-type has a different semantic value than its actual semantic value.

- (1a) The referent of ‘Everest’ as used by us at w_c differs slightly from its actual referent.
- (2a) The intension of ‘refers’ as used at w_c is the same as the intension of ‘refers’ as used at the actual world.¹³

If w_c is to satisfy (INT), then it must be that the intension of ‘refers’ as used by us at w_c maps the actual world to a relation that holds between ‘Everest’ and the referent of ‘Everest’ as used at w_c . But, given (2a), for this to be the case it must be that ‘Everest’ refers to the referent of ‘Everest’ as used at w_c . But, of course, ‘Everest’ refers to Everest. Since, by (1a), the referent of ‘Everest’ as used at w_c is not Everest, but some distinct, though largely overlapping, conglomeration of Himalayan rocks, it follows that (INT) is not satisfied. If there could be such a community, then there would be a proposition indistinguishable from the actual proposition expressed by (E) that was false.

‘Everest’ is a vague name. It has as its semantic value a large mass of rocks, but there are parts of the referent of this name—say pebbles at its base—which are not definitely so. According to Williamson’s epistemicism there are nearby possible worlds in which our use of language is such that ‘Everest’ refers to a sum of rocks mostly overlapping with the actual referent of ‘Everest’ but not including those as parts. The epistemicist will certainly allow, then, that there are worlds in which (1a) holds.

The claim to be defended, then, is that at least one of these worlds is one in which the intension of ‘refers’ is the same as its actual intension. It seems to me quite plausible that there should be such a world. One would have reason to reject this possibility if one were committed to a deep holism about semantic properties which held that any change in the semantic value of one term would bring about a change in the semantic value of every other term in the language. But this is completely implausible and would presumably be rejected even by those who think that semantic properties of words are determined by facts about language use as a whole. But, other than a commitment to this absurdly strong form of holism, I cannot see any principled reason to reject the existence of a world such as w_c . If we accept this possibility then we should accept (b).

Argument 2 We can argue for (b) with weaker assumptions. We need not assume that there is some world satisfying both (1a) and (2a). Instead, let us simply consider the class of nearby worlds in which the referent of ‘Everest’ differs from its actual referent, i.e., the class of worlds of which (1a) holds. As noted, the epistemicist certainly allows for the existence of such worlds. In order to avoid (b) one must maintain that (INT) holds of every such world. What I will argue is that maintaining this requires that one accept some bizarre and seemingly quite unmotivated restrictions on the intension of ‘refers’ as used at these worlds.

¹³ See Hawthorne (2006) for an argument that the intension of ‘refers’ as used in *all* nearby possible worlds is the same as the actual intension of this term. Hawthorne uses this to argue that the conditions imposed by (B*)—which he takes to be the conditions proposed by Williamson—are not, in fact, necessary for a case to be borderline. In what follows, I use weaker assumptions to argue that these conditions are not, in fact, sufficient.

Let w_a be some arbitrary world satisfying (1a). The first point to note is that if (INT) holds at w_a there will be an asymmetry between the intension of ‘refers’ as used at the actual world and the intension of ‘refers’ as used at w_a .

The intension of ‘refers’ as used at the actual world is sensitive to the slight differences in use between the actual world and w_a . In particular, this intension maps the actual world to a relation that holds between ‘Everest’ and Everest, while it maps w_a to a relation that holds between ‘Everest’ and a slightly different object, i.e., whatever is the referent of ‘Everest’ at w_a .

Assuming, however, that (INT) holds at w_a , it follows that ‘refers’ as used at w_a must lack such sensitivity. To see this first note that, since w_a is a nearby possible world in which our use of language is indiscriminably different from our actual use, everyone will go around confidently asserting the sentence “ ‘Everest’ refers to Everest”. From the fact that we know that ‘Everest’ refers to Everest, it follows, by epistemicist lights, that this sentence is true as uttered by us at w_a . For, according to the epistemicist, if we are to be safe in a judgment, then a judgment made using the same words in a nearby possible world must be correct. The intension of ‘refers’ as used at w_a , then, maps w_a to a relation which holds between the word ‘Everest’ and whatever is, in fact, the referent of ‘Everest’ as used at w_a . Assuming that (INT) holds, however, it follows that the intension of ‘refers’ as used at w_a will also map *the actual world* to a relation that holds between the word ‘Everest’ and whatever is, in fact, the referent of ‘Everest’ as used at w_a .

So, while the intension of ‘refers’ as used at the actual world is sensitive to the differences between our use of the expression ‘Everest’ in the actual world and in w_a , on the assumption that (INT) holds at w_a it follows that the intension of ‘refers’ as used at w_a must lack such sensitivity.

The claim that this asymmetry should be present at *every* nearby world in which the referent of ‘Everest’ differs from the actual referent should, I suggest, strike you as very strange. Why should any slight difference in our linguistic dispositions which would result in ‘Everest’ having a different referent necessitate this change in the behavior of the intension of ‘refers’? What is special about our actual dispositions which allows the intension of ‘refers’ be sensitive in this manner? I can think of no answers to these questions.

The claim that in every nearby possible world in which ‘Everest’ has a different referent the intension of ‘refers’ will lack this particular sensitivity to differences in use that is possessed by the intension of ‘refers’ as used at the actual world seems like an unmotivated hypothesis. But this is what is required if one is to maintain that (INT) holds for every such possible world.

The absurdity of this claim is even clearer when we consider that at every world w_a satisfying (1a) we will also assert sentences such as “Although ‘Everest’ refers to Everest, it might easily not have.” If we don’t want to undermine our own knowledge of the claim that we express using this sentence, then we cannot let our counterfactual selves be in error. But if we don’t attribute error here, then it must be that the intension of ‘refers’ as used at w_a is not such that it maps every nearby possible world to a relation holding between ‘Everest’ and the referent of ‘Everest’ as used at w_a . But then the claim that the intension must map the actual world *in particular* to such a relation seems completely without motivation.

I think that these reflections make it quite plausible that there are propositions that are false at the actual world and are expressed by (E) in nearby counterfactual situations in which our language use differs only slightly.

If the forgoing is right, then, given (B*), it follows that it is borderline whether ‘Everest’ refers to Everest despite the fact that we know that ‘Everest’ refers to Everest. This shows, I take it, that the proposed analysis is inadequate. What we wanted was to provide an account of what it is to be a borderline case using the resources of semantic indiscriminability that could explain why it is that it being borderline whether ϕ precludes knowledge that ϕ . (B*) fails to do this.¹⁴

6 Conclusion

Vagueness provides a barrier to knowledge. So does semantic indiscriminability. An appealing thought is that we can understand the former fact in terms of the latter. It is difficult, however, to make good on this idea.

For there to be a barrier to knowledge due to semantic indiscriminability is for there to be at least two worlds w_1 and w_2 suitably related to the actual world such that the proposition expressed by ϕ at w_1 is true at w_1 and the proposition expressed by ϕ at w_2 is false at w_2 . According to our first analysis, for it to be borderline whether ϕ just is for there to be a barrier to knowledge whether ϕ as a result of semantic indiscriminability. I noted two problems with this identification. (i) It leaves us without any obvious resources to say in virtue of what it is borderline whether ϕ at a world w when there are no worlds to which w stands in the same relation as the actual world stands to the worlds in $W_{@}$. (ii) There would seem to be a fundamental mismatch between the facts in virtue of which semantic indiscriminability precludes us from knowing whether ϕ and the facts in virtue of which it is borderline whether ϕ . The former are a matter of how things are at nearby possible worlds, while the latter are a matter of how things are at the actual world. Given these problems I argued that we should not *identify* it being borderline whether ϕ with there being a barrier to knowledge whether ϕ resulting from semantic indiscriminability.

The question, then, is whether one can provide an account of borderline cases in terms of facts about semantic indiscriminability that avoids these defects and is able to provide an explanation for why it being borderline whether ϕ is a barrier to knowledge that ϕ . I considered what I think is the most natural way of altering the first analysis that avoids the pitfalls of (i) and (ii). This account, however, failed to provide a general explanation for why borderline cases are barriers to knowledge. And this is unsurprising. For in order to resolve (i) and (ii) we were forced to ignore the facts about semantic indiscriminability that explain our ignorance, viz., the truth-values of indiscriminable propositions at nearby worlds.

¹⁴ Note that even if one were inclined to give up the claim that knowledge that ϕ is precluded by it being borderline whether ϕ (see e.g., Dorr (2003)), one may at least allow that it being borderline whether ϕ precludes it being definite that one knows that ϕ . But it seems that we definitely know that ‘Everest’ refers to Everest, and so one who is attracted to this weaker constraint should also find (B*) inadequate.

Perhaps there is some other way of spelling out the relationship between semantic indiscriminability and vagueness that can do better. But, given what has been said here, there seems to be good reason to think that the two phenomena are distinct. At the very least, until an adequate account of their relation has been offered we should be skeptical of Williamson's claim that vagueness consists in semantic indiscriminability.

Acknowledgments Thanks to Fabrizio Cariani, Cian Dorr, John MacFarlane, Mike Martin and Mike Rieppel for helpful comments on an earlier version of this paper.

References

- Dorr, C. (2003). Vagueness and ignorance. *Philosophical Perspectives*, 17, 83–114.
- Fine, K. (1975). Vagueness, truth, and logic. *Synthese*, 30, 265–300.
- Hawthorne, J. (2006). Epistemicism and semantic plasticity. In *Metaphysical essays* (pp. 185–210). Oxford: Oxford University Press.
- Kearnes, S., & Magidor, O. (2008). Epistemicism about vagueness and meta-linguistic safety. *Philosophical Perspectives*, 22, 277–304.
- Keefe, R. (2000). *Theories of vagueness*. Cambridge, MA: Cambridge University Press.
- Williamson, T. (1994). *Vagueness*. London: Routledge.
- Williamson, T. (2003). Vagueness in reality. In M. Loux & D. Zimmerman (Eds.), *Oxford handbook of metaphysics* (pp. 690–716). Oxford: Oxford University Press.
- Williamson, T. (2004). Reply to Mcgee and McLaughlin. *Linguistics and Philosophy*, 27, 113–122.