The Epistemicist Solution to the Sorites Paradox

Penultimate draft of chapter forthcoming in E. Zardini and S. Oms (eds.), *The Sorites Paradox*

(please cite published version)

Ofra Magidor, University of Oxford

§1. The epistemicist solution to the Sorites Paradox

Consider a typical instance of the Sorites Paradox[1]:

> A person who is 1000 millimetres (i.e. one meter) in height is not tall.
>
> $\forall n$(if a person who is *n* millimetres height is not tall, then a person who is n+1 millimetres in height is not tall).

Therefore,

> A person who is 2000 millimetres (i.e two meters) in height is not tall.

According to epistemicism, resolving the paradox requires no revision whatsoever of classical logic or semantics. The solution to the paradox is simply that the second premise is outright false.

In some ways, this is the most straightforward response to the paradox. First, it is entirely orthodox with respect to classical logic and semantics. Second, given classical logic and semantics the argument is valid and on the (practically uncontroversial) assumption that it has a false conclusion, it follows that it must have at least one false premise.[2] And while both premises have intuitive appeal, few would deny that if forced to choose, denying the second looks far more palatable than denying the first.[3]

However, reflection on this apparently straightforward response quickly leads to difficulties. After all, if the second premise is false, then its negation is true. This in turn means that there

[1] Two complications which I ignore: first, 'tall' is context sensitive. Second, what a person's height is might also be vague (for example do we include their hair or not?). These complications can be side-stepped by phrasing the paradox against a particular context and with a precise characterisation of the underlying height facts.

[2] Note that classical logic and semantics play a role mostly in establishing this last inference –most views of vagueness that reject classical logic or semantics nevertheless take the argument to be valid. However, see Magidor (2012) for an exception.

[3] For a dissenting voice see Enoch (2007).

is some specific number *n*, such that a person who is *n* millimetres in height is not tall, and yet a person whose height is just one millimetre greater is tall. Even worse, if we are to fully retain classical mathematics, then there must be a *sharp cut-off point* for 'tall': a precise real number *k* such that anyone who is less than *k* meters tall is not tall, and anyone who is over *k* meters tall is tall.[4] But it seems very hard to believe that words like 'tall' have such sharp cut-off points.

Epistemicists do not deny the existence of such sharp cut-off points. Rather, they add a second component to the view which is intended to make this surprising consequence more palatable. According to epistemicism, even though 'tall' has a sharp cut-off point, we do not and cannot know what this point is.[5]And moreover speakers mistakenly take their principled ignorance regarding the cut-off point as evidence that no such cut-off point exists.

The epistemicist solution to the paradox thus consists of two components. First, vague terms, just as non-vague ones, have ordinary classical semantic-values and do not require any revision of classical logic or semantics. In particular, each person is either tall or not tall, and for each person, it is either true that they are tall or false that they are tall. Second, although words like 'tall' have sharp boundaries, we do not and cannot know what these boundaries are, and this explains why we might be tempted to (falsely) conclude that such boundaries do not exist.

It is easy to see the advantages of epistemicism. It allows us to block the Sorites argument without requiring any revision of classical logic or semantics. Moreover, the view offers a straightforward account of higher-order vagueness, one that is entirely uniform with the treatment of first-order vagueness: the predicate 'definitely tall' also has a sharp cut-off point, although we do not and cannot know where it is. However the view also faces some significant challenges. Here are what I take to be the four main challenges for the view:

---

[4] This follows from the *completeness of the real numbers*, which says that any non-empty set of real numbers that has an upper bound has a least upper bound. Since the set {*l*: a person who is *l* millimetres in height is not tall} is non-empty and bounded it has a least upper bound *k*, which will serve as the cut-off point for 'tall'. Note that this leaves open whether or not a person who is exactly *k* millimetres in height is tall.

[5] When epistemicists say that we *cannot* know the sharp cut-off points of vague terms, they exclude special means of coming to know, such as testimony of an omniscient being (see Williamson (1997a): 926).

Challenge 1: Counter-intuitiveness

The view is highly counter-intuitive – it just seems very hard to believe that there is a specific real number which constitutes the boundary between being tall and not tall.

Challenge 2: the determination of meaning

It is commonly accepted that words have particular meanings in virtue of the way speakers use them. And yet, it does not seem like anything about the way we use the word 'tall' is sufficient to determine a particular sharp boundary.

Admittedly, according to semantic externalism it is not *just* the way we use words that determines their meaning: facts about the environment around us, as well as facts about which properties are more natural can play a role as well.[6] Thus, for example, whether speakers use the word 'water' in an environment where the water-like substance in their vicinity is H_2O or XYZ can determine which of the two substances the word 'water' picks out, and the fact that H_2O is a natural kind can explain why 'is water' in our mouth refers to the property of being H_2O rather than to the less natural property of being a clear tasty liquid. However, it is hard to see how this point helps with determining sharp boundaries for 'tall': the property of being over 1.781 meters in height exists just as much as the property of being over 1.782 meters in height and both properties seem equally (un)natural. Epistemicism thus faces the challenge of accounting for how vague words get assigned their sharp meaning.

Challenge 3: the explanation of ignorance

According to epistemicism we do not and cannot know the sharp cut-off point for 'tall'. But if 'tall' has a sharp cut-off point, why is it that we cannot discover it?

Challenge 4: the characterisation of vagueness

Vagueness seems to be a distinctiveness phenomenon: the words 'tall', 'bald', and 'intelligent' are vague, while 'even', 'prime', or 'exactly two meters long' are precise. How do we account for this difference? One possibility is to maintain that vague terms have a special kind of semantic-value, for example an incomplete semantic-value which gives rise to truth-value gaps. But this way of characterising vagueness is clearly not available to the

[6] See Putnam (1975) and Lewis (1983).

epistemicist. Nor can epistemicists merely appeal to the claim that we do not know precisely which people fall under the extension of ‘tall’. After all, ‘largest twin prime’ is a perfectly precise predicate, but we do not (and possibly cannot) know which number, if any, falls under its extension.[7] Epistemicism thus faces the challenge of accounting for this difference.

By far the most comprehensive defence of epistemicism is offered by Timothy Williamson, primarily in his book *Vagueness* (Williamson (1994)). In **§2**, I summarise Williamson's account and his response to the four challenges raised above. In **§3**, I discuss some of the objections that has been raised against Williamson's account and his way of addressing the four challenges. In **§4**, I offer concluding remarks.

§2. Williamson's Epistemicism

In this section I will focus on what has been the main defence of epistemicism: Timothy Williamson's. The key feature of Williamson's account consists in how he responds to the third challenge (the explanation of ignorance), but let us start with a brief discussion of Williamson's responses to the first two challenges.

Williamson's response to the first challenge (counter-intuitiveness) follows the standard epistemicist strategy: we are tempted to think that vague terms do not have sharp boundaries because, as much as we try, we cannot locate these boundaries or imagine how we might come to find them. However, once we realise that there is a principled reason for why we cannot locate these boundaries, the fact that we cannot find these sharp boundaries no longer provides evidence that they do not exist.[8]

What about the second challenge (the determination of meaning)? Williamson maintains that the meaning of vague words, just like precise ones, supervenes on their use (broadly construed to include facts about one's environment). This means that in any two possible worlds where the word ‘tall’ is used in precisely the same way, the word has exactly the same sharp cut-off point. However, he stresses that this metaphysical supervenience thesis does not

[7] A twin prime is a prime number x such that either $x+2$ or $x-2$ 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 whether there is a largest twin prime and if so which number it is.

[8] Of course, this in itself doesn't explain why there are other cases where we are ignorant (possibly principally ignorant) regarding some precise question and yet are not tempted to conclude that there is no fact of the matter. But I take it that Williamson isn't intending to offer here a full prediction of when speakers do or do not make such mistakes. Rather, he is trying to show how the mistake that (according to epistemicism) speakers are making is understandable (even more so given that the rather complex explanation for why we are principally ignorant of sharp-cut off points is unlikely to have occurred to speakers).

mean that there is any simple recipe that connects the use facts of a vague words to their meaning.[9]

Indeed, Williamson maintains that vague words possess a feature that makes the connection between their use facts and meaning especially fragile. Vague words are *semantically plastic*: their meaning is highly sensitive to the way they are used, so that even slight variations in their use patterns often make for a variation in their meaning. By way of illustration: if in the actual world 'tall' (for a person) picks out the property of being over 1.781 meters in height, then in a close possible world in which the use facts for 'tall' differs from ours only slightly, 'tall' picks out a different property (e.g. the property of being over 1.782 meters). This kind of semantic plasticity sits well with the observation that the property of being over 1.781 meters in height seems just as (un)natural as the property of being over 1.782, because the lack of natural candidates for the meaning of 'tall' can partially explain the semantic plasticity of the term. And it also sits well with Williamson's contention that while the meaning of vague terms supervenes on their use, the function mapping use patterns to meaning patterns is especially chaotic and intractable. The intractability of the supervenience function is important because it appears that even if we were to be given a very detailed description of all the use facts for 'tall', we still would not be able to say what the sharp boundaries of 'tall' are. It is thus important for Williamson that, even though the meaning of vague terms supervenes on their use, we do not (and cannot) know the precise details of the relevant supervenience function.

What about Williamson's response to the third challenge (the explanation of ignorance)? One might be tempted to assume that the intractable nature of the supervenience function is sufficient to account for our ignorance: since we do not know precisely how the meanings of vague terms supervene on their use, we cannot calculate their sharp cut-off points from their use facts. However, this explanation is unsatisfactory for two related reasons. First, this would only explain why we cannot know the sharp cut-off points of vague terms by calculating them from the use facts – but it does not explain why we do not know them using

[9] Indeed, a bit of reflection shows that there isn't a simple recipe even in the case of perfectly precise words. Consider for example 'even' (in the mathematical sense). While this term can be defined in a mathematically precise way, it was probably first used without an explicit stipulation of its meaning, which raises difficult questions regarding how many and what kinds of uses were required before the word acquired its meaning. Moreover, the Pythagoreans apparently had a theory of parity according to which the number one was neither 'even' nor 'odd' – whether this means that they had a false theory of parity or that their parity-like terms had slightly different meanings than ours, this case shows that the determination of precise meaning by use facts can be a complicated matter. (Cf. Williamson (1994), p. 207).

other, more direct means. Second, Williamson acknowledges that the mere fact that the supervenience function is highly complex, doesn't entail that it is principally unknowable. After all, there are many highly complex facts of science (including facts relevant to supervenience relations) which we do not necessarily take ourselves to be principally ignorant of. Thus, for example, you might think that the mental supervenes on the physical in very complex ways and still hold hope that science might be able to discover what the relevant correlation between the two domains is.

However, Williamson maintains that there is a crucial disanalogy between the supervenience of the mental on the physical and the supervenience of meaning on use. In the former case, scientific progress begins by amassing some direct observations concerning each of the two realms (e.g. preforming MRIs for the physical realm, and asking subjects whether they are in pain for the mental realm), and then forming scientific theories connecting the two domains. However, he maintains that in the case of vague terms we simply do not have sufficient direct access to facts about meanings. As he puts it: "Everyone with physical measurements *m* is thin' cannot be known *a posteriori* in a parallel way, for no route to independent knowledge of someone with physical measurements *m* that he is thin corresponds to asking someone whether he is in pain".[10] The upshot is that our ignorance of the supervenience function is derivative on our ignorance of the sharp cut-off points of vague terms, and thus Williamson wishes to offer a more direct explanation of the latter.

Williamson's direct explanation of our ignorance crucially relies (again) on the claim that vague terms are semantically plastic. An attractive principle in epistemology is that a necessary condition for a belief to count as knowledge is that the belief be *safe*. As a first approximation we can say that X's belief that p is safe just in case there is no close world in which X believes that p, but it is false that p. Suppose for example that you look at a jar full of marbles and form the belief that the jar contains exactly 135 marbles. Suppose that your belief was formed by pure guessing. Even if you happened to guess the correct number of marbles, most people would judge that you do not *know* that the jar contains exactly 135 marbles. One way to explain why your true belief does not constitute knowledge in this case is by pointing out that your belief isn't safe: there are close worlds where you make the same guess (and thus also form the believe that the jar contains 135 marbles), but where the jar contains 136 marbles, rendering your belief false.

[10] Williamson (1994): 204. See also the discussion in Kearns and Magidor (2008): 279 and §3.2 of this chapter.

Now consider the case of ‘tall’. Suppose that you form the belief that the sharp cut-off point for ‘tall’ is precisely 1.781 meters, and let us assume that your belief happens to be true. Still, claims Williamson, there are close worlds in which ‘tall’ is used slightly differently, and hence due to the plasticity of ‘tall’, the word has a slightly different meaning (e.g. it has a cut-off of 1.782). Moreover, since you are not sensitive to such slight variations in use (nor for that matter, to how exactly they affect the meaning of the term), then presumably in some such worlds you nevertheless believe that the cut-off for ‘tall’ is 1.781. But this latter belief would be false, rendering your actual belief unsafe. The upshot is that due to semantic plasticity, even if you were to form a true belief concerning the cut-off point for ‘tall’, that belief would not constitute knowledge.

So far we have focused on our ignorance of *semantic cut-off claims* – claims such as ‘The word ‘tall’ picks out the property of being more than 1.781 meters in height’. However, it is important to observe that we also have a corresponding ignorance which is not directly about the word ‘tall’ or its semantic-value: we are also ignorant about the boundaries of the property of tallness. Assume for example that the cut-off for ‘tall’ is 1.781 meters. We do not know claims such as the following: ‘A person is tall if and only if they are over 1.781 meters in height’ (call this a *non-semantic cut-off claim*). Interestingly, it is such non-semantic claims that appear in typical versions of the Sorites Paradox. The problem, however, is that since these claims are not about the word ‘tall’, then on the face of it, semantic plasticity seems irrelevant to explaining our ignorance of these. Relatedly, note that if it is true that a person is tall if and only if they are over 1.781 in height, then it is *necessarily* true.[11] Of course, in other worlds the sentence expressing this necessary truth might express a different proposition which is false, but that is irrelevant to the necessity of the truth actually expressed. (Compare: it is necessarily true that two plus two equals four, even though in some other possible worlds, the word ‘four’ denotes the number five, and hence the sentence ‘two plus two equals four’ expresses a falsehood). But if the non-semantic cut-off claim is necessarily true, any belief in it trivially satisfies safety: since it is not false in any world, one does not falsely believe it in any close world.

Note, however, that this is an instance of the more general challenge of how to apply the safety principle to necessary truths. Suppose for example that you form the belief that

[11] This is obvious if one thinks that the property of tallness just is the property of being 1.781 in height, but even those who have a fine-grained conception of properties that distinguishes between the two, would presumably want to accept that whether or not a person is tall supervenes on their height.

256x31==7936 by simply guessing the result of the multiplication. Just as with the case of the jar of marbles, your belief does not constitute knowledge. However, this time there is no close world in which you falsely believe that 256x31==7936 simply because there is no possible world (close or otherwise) where this proposition is false. In light of such examples, defenders of safety principles usually extend the principle so that it applies more liberally. According to *extended safety*, your belief that p is safe just in case there is no close world in which you have a *similar* belief that q, and it is false that q. For example in the above case there is a close world *w* in which, when evaluating 256x31, you make a different guess – say 7937. In that world you believe that 256x31=7937, and that belief is false. Although your belief in *w* has a different content than your actual belief, it is sufficiently similar to your actual belief for the latter to violate extended safety, entailing that it does not constitute knowledge.

Williamson utilises this more liberal conception of safety to argue that semantic plasticity leads to ignorance of non-semantic cut-off claims as well. Suppose you form a true belief in the claim that a person is tall if and only if they are over 1.781 meters in height. On Williamson's account, there is a close world *w*, one with subtly different use facts, in which you form a very similar belief, one that you would express in *w* by using exactly the same words ('A person is tall if and only if they are over 1.781 meters in height'). However, due to the semantic plasticity of 'tall', in *w*, the term picks out a different property (say, the property of being over 1.782 meters) and thus your belief in *w* is false. Of course, your *w*-belief has a slightly different content than your actual belief, but Williamson maintains that the close linguistic similarity between the way you phrase these two beliefs (you use the same words, which have nearly identical meanings), is sufficient to render them similar for the purposes of extended safety, thus ensuring that your actual belief does not constitute knowledge.

The explanation of ignorance in terms of semantic plasticity also plays a role in the way Williamson responds to the fourth challenge (the distinctiveness of vagueness): our ignorance in cases of vagueness is ignorance due to a distinct source – namely, semantic plasticity. While we may be ignorant (perhaps irredeemably ignorant) of which number if any fall under the extension of 'largest twin prime', our ignorance in the latter case clearly has nothing to do with plasticity.

In addition to addressing the four challenges, Williamson's account has the following attractive feature: it predicts that while we cannot know truths of the form ''tall' refers to the property of being over *k* meters in height', we *can* know truths of the form ''tall' refers to the property of tallness'. For consider your belief that 'tall' refers to the property of tallness. Due to semantic plasticity, there is a close world *w* in which 'tall' refers to a slightly different property (call it *tallness** - which will be the property of being over *k** meters in height). As above, there is a close world *w* where you form a belief that you would express by saying ''tall' refers to the property of tallness''. But while this belief has a slightly different content than your actual belief (in *w* it expresses the content that 'tall' refers to tallness*), this time your *w*-belief is *true* in *w*, and thus does not render your actual belief unsafe. This is an attractive result, not only because it is intuitive that we do know that 'tall' refers to tallness, but also because it shows that there is a sense on which we know what 'tall' means - despite not knowing the sharp cut-off point of the term. (Compare: we may not be able to describe the extension of 'twin prime' in an informative way, but we nevertheless know what the term means).

§3 Objections to Williamson's account

Williamson offers a powerful defence of epistemicism, but a range of objections have been raised against his account. For many philosophers, the very idea that words like 'tall' have sharp cut-off points seem too far-fetched to accept - to use Lewis's picturesque language, epistemicism is often met with an 'incredulous stare'. In this section, however, I would like to focus on more concrete objections to the account. I discuss in turn objections to Williamson's response to each of the four challenges.

§3.1 counter-intuitiveness

Recall that Williamson argues that one thing that explains why speakers believe that vague predicates do not have sharp boundaries is that they mistakenly infer from their inability to locate these boundaries that no such boundaries exist. Suppose that *k* is the correct sharp cut-off point for 'tall'. Williamson's account focuses on explaining why speakers do not *know* that *k* is the correct cut-off by explaining why even if we truly believed that it is the cut-off point, that belief would not constitute knowledge. However, Keefe points out that there is a more basic fact that needs to be explained: namely, why speakers do not even *believe* that *k* is the sharp cut-off point.[12] Moreover, it seems that it is this latter fact which is relevant to addressing the counter-intuitiveness challenge: when we are trying to explain why speakers

[12] Keefe (2000): 71-72. Also relevant is Horwich (1997) who maintains that the very phenomenon of vagueness is characterised by an unwillingness to believe borderline statements rather than in the inability to know them, but see the response in Williamson (1997b): 945-6.

find it hard to believe that there are sharp cut-off points, what matters is that they do not *believe* (rather than do not know) of any particular number that it is a sharp cut-off point.

Williamson could respond by saying that speakers do not typically believe that *k* is the sharp cut-off point because they realise that such a belief would not constitute knowledge. But as Keefe points out, this leaves the question of why speakers think that such beliefs would not constitute knowledge. It obviously will not do to say that it is because speakers accept Williamson's explanation of ignorance – after all, the vast majority of speakers have never even considered this explanation.

Here the dialectic between Keefe and Williamson becomes rather delicate: on the one hand, it is clear that even if Williamson's explanation of ignorance is correct, it does not directly figure in the beliefs of ordinary speakers. On the other hand, even Keefe does not deny that ordinary speakers do not believe that they know (or can know) the sharp boundaries of vague terms.

Perhaps the following analogy can help make progress: suppose that most people believe that there is no fountain of youth. I argue that they are wrong, the fountain of youth does exist but it is impossible to find it, which is why everyone else falsely assumes that it doesn't exist. You ask me why the fountain is impossible to find. I respond that it is impossible to find because the Gods keep it hidden: every time someone comes close to discovering the fountain, the Gods immediately move it to another location. You press me by saying that this only explains why no one *knows* where the fountain of youth is, not why there aren't locations for which speakers *believe* (either falsely or as a luckily true guess) that the fountain is located, especially since other speakers haven't even considered my theory about the Gods moving the fountain. How should I respond to this latter complaint? It is clear that the reason speakers do not believe the fountain is located in any particular location does not involve their accepting my theory about the Gods – speakers don't think the fountain is located at *l* simply because they think it does not exist. However, this does not mean that my explanation of their ignorance is irrelevant to accounting for their beliefs. After all, if the Gods hadn't kept the fountain hidden, then perhaps people would have eventually found the fountain and thus altered their belief that it does not exist.[13] Thus at least in so far as the case of hidden sharp cut-off points is analogous to that of the hidden fountain of youth, it seems that

[13] And note that if they merely formed a true belief that isn't knowledge concerning the claim that the fountain exists, their belief would arguably be less robust (cf. Williamson's discussion of the burglar who merely has a true belief rather than knowledge that there is a diamond in the house in Williamson (2000): 62-64).

Williamson's explanation of ignorance is after all relevant to explaining the lack of belief in sharp cut-off points, even if speakers are not aware of his account.

§3.2 The determination of meaning

Williamson's response to the challenge concerning the determination of meaning is to maintain that the meaning of vague terms supervenes on their use, although the function connecting use facts to meaning facts is chaotic and intractable. Keefe notes that merely asserting that meaning supervenes on use does not help appease the intuition that in cases of vagueness, meaning is underdetermined: it at least *seems* as if nothing determines whether by 'tall' we pick out the property of being over 1.78213 meters rather than over 1.78214 meters in height.[14]

Another worry with the claim that the meaning of vague terms supervenes on their use concerns the question (which I have already discussed briefly above) of why we cannot come to know the relevant supervenience function: we do not (and apparently cannot) know facts of the form 'If speakers were to use language thus and so, the cut-off for 'tall' would be precisely *k*'.[15]

As noted above, the mere fact that the supervenience function is highly complex is insufficient to explain our ignorance, because many other scientific claims - which we do not take ourselves to be principally ignorant of - are highly complex as well. As we have seen, Williamson's explanation for why the meaning-on-use supervenience function is unknowable relies on the claim that we do not have direct access to a sufficient range of meaning facts, and hence cannot come to know how the supervenience function works using standard scientific methods. But it is not clear how satisfying this response is: after all, there are many facts about the meanings of vague terms that we *do* have direct access to (for example that someone who is 1.50 meters in height is not tall or penumbral connections involving vague terms). Just as forming a scientific theory concerning the correlation between the mental and the physical does not require an observation of every single brain in the world, one might maintain that the meaning facts we do have direct access to could provide us with a sufficient basis to gain scientific knowledge of the relevant supervenience function.[16]

[14] Keefe (2000): 77.
[15] Note that an explanation in terms of semantic plasticity would not work here, as this sentence expresses a truth in all close worlds.
[16] See Kearns and Magidor (2012): 279, 298-299 for more details.

In light of both of these challenges, one might propose a different variant of epistemicism, one which shares with Williamson the view that vague terms have sharp cut-off points, but relies on an unorthodox view according to which meaning does not (or at least not fully) supervene on use or other non-semantic facts.[17] Suppose our use facts leave open whether 'tall' picks out the property of being over 1.781 meters or the property of being over 1.782 meters in height. According to this proposal, for each possible world containing the same use facts as ours, there is a primitive semantic fact which states which of the two properties 'tall' picks out. Of course, this proposal might be thought to be even more counterintuitive and radical than Williamson's: it not only postulates sharp cut-off points, but also primitive semantic facts. However, Kearns and Magidor (2012) argue that there are reasons completely independent of vagueness to accept the view the meaning doesn't supervene on use. For example, they show that similar considerations to those that have driven theorists to reject the supervenience of the mental on the physical (ones involving worlds with purely mental beings or phenomenal zombies) can generalise into arguments against the supervenience of the semantic on the non-semantic.

This proposal is well-equipped to address the above objection concerning the determination of meaning: our intuition that the use of vague terms does not fully determine a precise cut-off point is actually correct, and the reason we cannot come to know the relevant supervenience function is that no such function exists. It should be noted, however, that this alternative form of epistemicism faces other challenges. For example, while this view does not take meaning facts to fully supervene on use facts, presumably it will still accept that meaning facts are *constrained* by use facts. But this leaves open the question of how use facts constrain the meaning facts, a question which is also relevant to the issue of how, on this view, to account for higher-order vagueness.

§3.3 The explanation of ignorance

Suppose in the actual world you form a belief that you express by saying 'A person is tall if and only if they are over 1.781 meters in height', and let us suppose that this belief is true. Recall that Williamson argues that this true belief does not constitute knowledge because there is a close world *w* in which the following hold:

(i) *Plasticity*: The use facts in *w* differ slightly from ours, entailing that the sharp cut-off for 'tall' in *w* is slightly different.

[17] See Kearns and Magidor (2012) for a defence of this view.

(ii) *Similar Belief*: In *w*, you nevertheless form a very similar belief to your actual one, one that in *w* you would also express by using the same words: 'A person is tall if and only if they are over 1.781 meters in height'. (Of course, given (i) this belief is false in *w*).

Moreover,

(iii) *Meta-linguistic Safety:* The linguistically-similar false belief you have in (ii) entails that your actual belief is not safe and thus does not constitute knowledge.

Objections have been raised against each of the three components of this explanation.

Let us begin with *Semantic Plasticity*: why should we think that in close worlds where 'tall' is used slightly differently, it has a different meaning? Note for a start that we must accept that in at least some cases, two worlds which differ only slightly in use facts, differ in meaning facts. Indeed this point has nothing to do with vagueness. Consider a precise word such as 'prime'. It is not a necessary truth that this particular combination of graphic symbols or sounds picks out the property of being prime - there is some possible world w_n in which speakers use the word so differently than we do, that it refers to some completely different property (e.g. the property of being a cat). Moreover, we can construct a series of worlds $w_1,\ldots,w_n$ connecting the actual world w_1 to w_n such that in any two adjacent worlds of the series, the use facts differ only slightly. Since in the actual world 'prime' picks-out the property of being prime and in w_n it does not, there must be two adjacent worlds in the series where the meaning of 'prime' shifts, entailing that there are cases where small shift in use makes for a shift in meaning.

However, for precise terms like 'prime', we can assume that such shifts between close worlds are very rare. One plausible hypothesis of why this is so involves Lewis's theory of eligibility.[18] According to this theory, the meaning of a term is determined by a combination of two factors: first, how well it fits with our use facts, and second how natural the property it picks out is. More natural properties serve as 'reference magnets' – properties that are easier to refer to, which means a wider range of uses succeed in latching on to the property (as long as it's a reasonably good fit for use). Since the property of being prime is a natural one, 'prime' is not semantically plastic: many different use patterns for 'prime' pick out the

[18] See Lewis (1983).

property of being prime. It is only when use facts diverge so significantly from our current use so as to render primeness no longer a reasonable fit for our use facts that the meaning of the term shifts. However, in all close worlds, ‘prime’ has the same meaning it actually does.

Why not think that ‘tall’ follows a similar pattern, namely that it picks out the same property in the actual world and in all close worlds? As we have seen, Williamson’s response is that this has to do with the lack of natural properties as candidate meanings for ‘tall’: since the property of being over 1.781 meters in height seems just as natural as the property of being over 1.782 meters, it is hard to see how a range of different use facts would consistently pick out one property over the other.

Although it *seems* plausible that all candidate meanings for ‘tall’ are equally natural, some authors have argued that perhaps naturalness is less transparent than is usually assumed in these discussions. That is, it might turn out that unbeknownst to us, the property of being over 1.7812 is much more natural than any other properties in the vicinity.[19] More cautiously, Mahtani argues that although it is implausible *tout court* that some such height property is natural, *conditional* on accepting the (in itself implausible) claim that vague predicates have sharp cut-off points, the additional claim that there are such hidden naturalness facts might be quite likely: after all, it would be far less surprising that a word like ‘tall’ refers to the property of, say, being over 1.7923 meters in height if it turned out that this property serves as a natural reference magnet. If these suggestions are correct, vague terms are not semantically plastic which would undermine Williamson’s explanation of ignorance.[20]

This kind of objection to semantic plasticity is weak in the sense it merely questions whether vague terms must be semantically plastic, but it does not give any reason to think that vague terms are *not* semantically plastic. However, a stronger direct argument against plasticity is suggested by Dorr and Hawthorne.[21] Suppose that George is telling us about his date with Bill, and mentions that his date wasn’t tall. In this situation we can truly utter counterfactuals such as the following:

> If George had gone out with Paul instead, he would have said his date was tall!

[19] See Mahtani (2004) and Cameron (2010).
[20] Note though that in that case one could offer an alternative explanation of our ignorance of sharp cut-off points in terms of our ignorance of naturalness facts (see Cameron (2010)).
[21] See Dorr & Hawthorne (2014), and also related remarks concerning temporal plasticity in Kearns & Magidor (2008): 286-7.

The problem is that if ‘tall’ is semantically plastic, it is hard for utterances of such counterfactuals to come out as true. For consider the closest possible world *w* in which George goes out on a date with Paul rather than Bill. In *w*, the use facts of ‘tall’ are presumably slightly different than the actual ones (for example, in *w* but not in the actual world, George utters the sentence ‘Wow, you are pretty tall!’ when he first sees his date, adding one more use of the word ‘tall’). So by semantic plasticity, ‘tall’ picks out a slightly different property (call this property ‘tall*’). But then in *w* George does not say that his date was tall, since in *w* no sentence George utters in the situation expresses the proposition that his date is tall.[22] Of course, in *w* George might well utter the words ‘My date was tall!’, but remember that in *w* those words amount to George’s saying that his date was tall*. Defenders of semantic plasticity thus at the very least face the challenge of explaining how such a counterfactual can be true.[23]

Next, consider *Similar Belief*, namely the claim that if in the actual world you have a belief that you would phrase by saying ‘A person is tall just in case they are over 1.781 meters in height’, then in close worlds you have a corresponding belief that you would put using exactly the same words. Why should we think that in close worlds you would still have a belief that you would express using exactly the same words? Suppose in the actual world you formed your belief about the boundaries of tallness by quick reflection or a guess concerning what the cut-off of tallness might be. In that case, it is very plausible that your guess is not sensitive to subtle differences in use facts, and hence in at least some close worlds where use facts differ slightly, you would endorse exactly the same sentence. The problem, however, is that we do not really need Williamson’s sophisticated explanation via semantic plasticity to tell us that guessing or crude reflection are not good methods for acquiring knowledge about subtle matters. What if instead you formed your belief via some very reliable careful method for discovering cut-off points? In that case, we can expect that in worlds where the cut-off point for ‘tall’ is, say, 1.782 (rather than 1.781) you would apply the same careful method and reach a conclusion that you would put using a different sentence (‘A person is tall just in case they are over **1.782** meters in height’). Of course, it is hard to imagine what kind of method could play this role. But Williamson’s explanation of ignorance was supposed to go beyond

[22] Of course if tallness is the property of being over, say, 1.798 meters in height, George *could* utter the sentence ‘My date was over 1.798 meters in height’, but we are assuming that in close worlds George does not make such odd reports about his date.
[23] Dorr and Hawthorne discuss several attempts to respond to this challenge, but show they face serious difficulties. On the other hand, they also raise problems for rejecting plasticity altogether.

merely noting that we cannot imagine a method by which we might discover the sharp cut-off points of vague terms.[24]

Finally, consider *Meta-Linguistic Safety*. Williamson's account relies on a particular version of the safety principle: one according to which your belief that *s* counts as unsafe (and hence does not constitute knowledge), if there is a close world in which you have a corresponding false belief which is *linguistically* similar, namely a belief that you would phrase using the same words. However, Kearns and Magidor (2008), argue that this kind of meta-linguistic safety is not a necessary condition on knowledge. To do so, they provide a series of cases where an agent has a belief that violates meta-linguistic safety, but arguably nevertheless constitutes knowledge.

Consider for example the following case: suppose that Joe and Bill are two identical twins. The use facts of the community determine that the name 'Joe' picks out the first twin Joe, and the name 'Bill' picks out the other twin, Bill. However, suppose that since it is very easy to confuse the two twins, many times when speakers see Bill they refer to him using the name 'Joe' and vice versa. Let us suppose that these confusions are so frequent, that had a few more speakers mixed the names up, this would have simply shifted the semantic-value of these names: thus, there is a close world *w* where the name 'Bill' refers to Joe and the name 'Joe' refers to Bill.

Now, suppose that Joe spends a year in Australia, far away from his twin brother Bill. In Australia, Joe becomes friends with Jill. Jill comes to know Joe by the (correct) name 'Joe', does not know he has a twin brother, and is under no risk of ever encountering his twin brother. After interacting a lot with Joe, Jill forms the belief that Joe is a nice guy. Suppose Jill's belief is true and based on a lot of excellent evidence (Joe has been a good friend to her, volunteers in the community, and so forth). It seems clear that Jill's belief constitutes knowledge. However, Jill's belief violates meta-linguistic safety. This is so, because in *w*, Jill still has a belief that she would phrase using the sentence 'Joe is a nice guy' (simply because she would mistakenly assume that 'Joe' in *w* picks out Joe). However, for reasons entirely unavailable to Jill, the name 'Joe' in her mouth in *w* picks out Bill. And assuming that Bill is not nice, her *w*-belief would be false, rendering her actual belief not meta-linguistically safe.

[24] Versions of this objection are raised in Keefe (2000): 74-75, Gomez-Torrente (2002): 112-13, and Kearns and Magidor (2008): 283-288. As Keefe points out, part of the problem is that since we do not really actually believe in particular cut-off points, it is hard to conjecture what methods we would have used if we had such beliefs.

It is hard to see how her unfortunate linguistic mistake in *w*, though, can impact Jill's knowledge about Joe's niceness in the actual world.

In response to this objection, Williamson could respond as follows: Meta-linguistic safety isn't, in full generality, a necessary condition on knowledge. Rather, knowledge requires extended safety, but which beliefs count as similar in the sense required for extended safety is a highly delicate matter that might not be susceptible to philosophical analysis (at least not to non-circular analysis, i.e. one that does not appeal to the concept of knowledge). And while in the particular case of vague terms it turns out that meta-linguistically similar beliefs count as similar, this might not generalise to other cases.[25]

Indeed, similar manoeuvres could be made in response to (nearly) all the objections raised above.[26] That is, the epistemicist could address the worries about whether vague terms really are semantically plastic by stipulating that they are, or whether speakers really do form beliefs using the same sentence in close worlds by stipulating that they do. The idea would then be that the various stipulations figuring in the explanation of ignorance are not intended to be motivated independently of their roles in this explanation. The explanation of ignorance can be thought of in the first instance as merely a model for how the ignorance of vague terms could in principle arise, where the reason for accepting this model as realistic being that it successfully explains why we do not know the sharp boundaries of vague terms.[27]

The question then turns on whether this more modest interpretation of the explanation of ignorance is dialectically satisfying. The answer depends in part on what one expects to achieve from such an explanation. If one is already convinced of the existence of sharp boundaries (for example, because it is entailed by classical logic and mathematics), one might find this sort of model for why they are unknowable highly illuminating. On the other hand, if one starts out as very sceptical that vague predicates even have sharp boundaries, but finds the existence of sharp boundaries easier to accept if coupled with a compelling independent

[25] For a general defence of the claim that safety can only be analysed using the concept of knowledge see Williamson (2000). For a critical discussion of the use of this claim in the context of explaining our ignorance of vague terms, see Kearns & Magidor (2008): 297 and Sennet (2012): 279.

[26] I say 'nearly' because I don't think this point helps to address Dorr & Hawthorne's worry.

[27] By analogy, consider an interpretation of Lewis's account of counterfactuals which accepts that there is no non-circular understanding of similarity amongst worlds (i.e. that similarity is explained using counterfactuals). Suppose one wanted to explain why, on one's view, counterfactuals are non-transitive: $A \Box\!\!\rightarrow B$ and $B \Box\!\!\rightarrow C$ does not entail $A \Box\!\!\rightarrow C$. The Lewisian account does offer an explanation in the sense that it provides a model for how this can come about (i.e. if the closest A world is a B&¬C world, but there is an even closer world where ¬A, B, and C). However, it will not provide any independent argument that this structure is in fact realised, and thus would not convince someone who is sceptical that there are any actual counterexamples to the transitivity of counterfactual conditionals.

argument to the extent that if such boundaries exist they must be principally undiscoverable, this more modest interpretation of Williamson's explanation of ignorance might well fall short.

§3.4 The distinctiveness of vagueness

The fourth challenge for epistemicism was to account for what is distinctive about vagueness. Williamson's response is to maintain vagueness is associated with a distinctive kind of ignorance – one arising due to semantic plasticity.

A general challenge to the characterisation of vagueness in terms of semantic plasticity is suggested by Sennet.[28] Imagine a possible world *w*, where a community of speakers speak a language very much like English, and indeed use 'tall' just as we do. However, in all worlds close to *w*, there is a daemon who ensures that 'tall' has exactly the same sharp cut-off point as in *w*, ensuring that the term is not semantically plastic. (One way the daemon can achieve this is by adding more use facts to exactly balance out variations in use among the other speakers in the community.) The problem is that it seems that in *w*, the community's use of 'tall' is just as vague as ours and they are equally ignorant of its sharp cut-off point, despite the fact that the term is not semantically plastic.

A natural response to this challenge is to maintain that even if in one sense of 'close' the daemon is present in all close worlds, that is not the sense relevant to our assessment of knowledge. On the relevant understanding of closeness, there are some close worlds on which the daemon is not present allowing the term to be plastic after all. Whether we accept this response is related to the issue discussed at the end of §3.3, as to whether the explanation of ignorance is intended to rely on an independent understanding of closeness.

Let me turn to an entirely different challenge to the characterisation of vagueness in terms of a particular kind of ignorance due to plasticity. To do so, we first need to note that not every sentence containing vagueness involves ignorance. The word 'tall' is vague, but we still know that 'A person who is two meters in height is tall' or that 'If *x* is taller than y, then if *x* is tall than so is *y*'. Rather, ignorance only arises in the case of *borderline* statements. Thus to be more precise, the hypothesis is that borderline statements are characterised in terms of the

[28] Sennet (2012)

distinctive kind of ignorance, and vagueness is characterised via its susceptibility to borderline cases. The question is then whether Williamson's account offers an adequate characterisation of borderlineness.

Let us suppose that we are generally sympathetic to Williamson's explanation of ignorance of particular borderline cases. Still, it is surprisingly difficult to extract from this explanation a set of necessary and sufficient conditions for a sentence's being borderline. Following Williamson's explanation of ignorance, we might try the following: a true sentence '*s*' is borderline[29] just in case there is a close possible world *w* in which:

(i) due to semantic plasticity, 's' expresses in *w* a different proposition than it actually does

(ii) the proposition '*s*' expresses in *w* is false in *w*

(iii) If in the actual world speakers have a belief that they would express using '*s*', they have a similar belief that they express using the same words in *w*.

As Caie (2012) and Magidor (*forthcoming*) argue, however, this explanation will not work for the following reason. Suppose that in the actual world the sharp cut-off point for 'tall' is 1.781 meters, and Jill is 1.8 meters tall (let us assume this is sufficient to make her definitely tall). Now consider a close world *w* in which, due to the semantic plasticity of 'tall', the sharp cut-off for 'tall' is 1.782. Assume also that in *w*, Jill is a bit shorter than she actually is: in *w* she is only 1.78 meters tall. The problem is that according to the above account 'Jill is tall' would come out as borderline: the sentence expresses a different proposition in *w* (namely, the proposition that Jill is taller than 1.782 meters in height) and that proposition is false in *w* (because Jill's height in *w* is 1.78 meters < 1.782 meters). Finally, we can simply stipulate that the doxastic attitude of speakers are such as to satisfy clause (iii) of the condition. However, as Jill is definitely tall in the scenario, the prediction that the sentence is borderline is incorrect.

It is clear what went wrong. When assessing whether 'Jill is tall' is borderline, we are interested in close worlds that differ only in the meaning of the word 'tall', not in the underlying conditions relative to which the truth of the sentence is assessed, such as Jill's height. However, it is not obvious how to translate this observation into a condition on borderlineness. Magidor (*forthcoming*) proposes that we can make progress by appealing to

[29] A false sentence 's' can then be said to be borderline iff its true negation is borderline.

the notion of metaphysical ground. The proposal is to require that whatever facts ultimately ground the fact that *s* (in the actual world) would be held fixed in the close world *w*. Thus, in the above example, what actually grounds the fact that Jill is tall is her precise height (i.e. the fact that she is 1.8 meters tall). But then we should only consider close worlds in which Jill's height is also 1.8 meter, and the above problem would not arise. Even if in a close world *w* 'tall' has a slightly different cut-off than it actually has, 1.8 meters would presumably be above that cut-off, and the proposition expressed by 'Jill is tall' in *w* would be *true* rather than false, so the account would predict that the sentence is not borderline, just as we expected.

However, while this addresses the problem with the current example, it may be too restrictive because of examples such as this. Consider the sentence 'The word 'frequent' is used frequently'. Suppose that 'frequent' is used fairly frequently, so that it is borderline whether the word is used frequently or not. The problem is that part of what metaphysically grounds the fact that 'frequently' is used frequently, is precisely the complete use facts for the word 'frequently'. But if we only consider close worlds in which we are holding fixed these use facts, then due to the supervenience of meaning on use, the word would have exactly the same sharp cut-off as it actually does, so clause (i) of the condition (the requirement that due to semantic plasticity the sentence expresses a different proposition in *w*) will not be satisfied, and the sentence would falsely be predicted to be definite rather than borderline. The problem is that we want to hold fixed the underlying conditions relative to which the sentence is assessed, but in doing so, we are also accidently holding fixed the meaning of the relevant vague words. It is thus not clear that there is any systematic way of *only* holding fixed the underlying conditions, and thus get an adequate characterisation of borderlineness in terms of the distinctive kind of plasticity.

§3.5 concluding remarks

The epistemicist view of vagueness, offers a straightforward response to the Sorites Paradox. As we have seen, though, the view faces several challenges. We have discussed Williamson's attempt to respond to these challenges and some objections to his responses.

Even if one finds these objections compelling, one should not be too quick to dismiss epistemicism in favour of alternative views. For one thing because in discussing these objections one should be careful not to hold epistemicism to higher standards than the competing views. While our discussion focused on the challenges as they apply to epistemicism, many of these challenges can be equally raised against competing views. Consider the determination of meaning: supervaluationists face the challenge of explaining how our use facts manage to pick out one set of admissible precisifications over another. Or consider the explanation of ignorance: fuzzy logicians face the challenge of why we do not (and apparently cannot) know exactly which degree of truth 'Jill is tall' receives.[30] And even concerning the distinctiveness of vagueness, non-classical views arguably have difficulties in giving a criterion that distinguishes vagueness from other apparent cases of truth-value gaps, e.g. ones due to semantic paradoxes such as the Liar.

Moreover, we should not lose sight of the significant advantages that epistemicism offers: both its full retention of classical logic and semantics, as well as its straightforward treatment of higher-order vagueness. These considerations give us strong reasons to keep developing the details of the view, rather than abandoning it all together.[31]

---

[30] Of course supervaluationists or fuzzy logicians might ultimately wish to deny that we do succeed in picking out a specific set of admissible precisifications or degrees of truth, but these ways of developing the views bring with them other challenges.

[31] Thanks to the editors of this volume for their helpful comments on this chapter. Thanks also to the Leverhulme Trust for their support.

References

Caie, M. (2012), ‘Vagueness and semantic indiscernibility’, *Philosophical Studies* 160: 365-377.

Cameron, R. (2010), ‘Vagueness and naturalness’, *Erkenntnis* 72: 281-293.

Dorr, C. and Hawthorne, J. (2014), ‘Semantic plasticity and speech reports, *Philosophical Review* 123: 281-338.

Enoch, D. (2007), ‘Epistemicism and nihilism about vagueness: what is the difference?’, *Philosophical Studies* 133: 285-311.

Gómez-Torrente, M. (2002), ‘Vagueness and margin for error principles’, *Philosophy and Phenomenological Research* 64: 107-125.

Horwich, P. (1997), ‘The nature of vagueness’, *Philosophy and Phenomenological Research* 57: 929-935.

Kearns, S. and Magidor, O. (2008), ‘Epistemicism about vagueness and meta-linguistic safety', *Philosophical Perspectives* 22: 277-304.

Kearns, S. and Magidor, O. (2012), ‘Semantic sovereignty’, *Philosophy and Phenomenological Research* 58: 322-350.

Keefe, R. (2000), *Theories of Vagueness*, Cambridge: Cambridge University Press.

Lewis, D. (1983), ‘New work for a theory of universals’, *Australasian Journal of Philosophy* 61: 343-377.

Magidor, O. (2012), ‘Strict finitism and the happy sorites’, *Journal of Philosophical Logic* 41: 471-491.

Magidor, O. (*forthcoming*), ‘Epistemicism, distribution, and the argument from vagueness’, forthcoming in *Nous*.

Mahtani, A. (2004), ‘The instability of vague terms’, *Philosophical Quarterly* 54: 570-576.

Putnam, H. (1975), ‘The meaning of ‘meaning’’, *Minnesota Studies in the Philosophy of Science* 7:131-193.

Sennet, A. (2012), ‘Semantic plasticity and epistemicism’, *Philosophical Studies* 161: 273-285

Williamson, T. (1994), *Vagueness*, London: Routledge.

Williamson, T. (1997a), ‘Précis of vagueness’, *Philosophy and Phenomenological Research* 57: 921–928.

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

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