Epistemicism without metalinguistic safety

Justin Khoo

In preparation for: *Linguistic Luck*, Abrol Fairweather and Carlos Montemayor (eds.)

Abstract

Epistemicists claim that vague predicates have precise but unknowable cutoffs. I argue against against the standard, Williamsonian, answer, that appeals to metalinguistic safety: we can know that p even if our true belief that p is metalinguistically lucky. I then propose that epistemicists should be diagonalized epistemicists and show how this alternative formulation of the view avoids the challenge. However, in an M. Night Shyamalan-style twist, I then argue we should not be diagonalized epistemicists either.

Terry has volunteered to be zapped by the Shrink Ray 3000, a device that causes the target to shrink in height quickly and continuously over one minute. We type in the setting that causes a rate of diminishment of one quarter inch per second. Terry's height at t_0, just before he is zapped, is 6'7", which is sufficient for him to be tall (for an average American 30 year old man). After sixty seconds of zapping, at t_{60}, his height is 5'4", which is sufficient for him to be not tall (for an average American 30 year old man). So, at some point in the process, he went from being tall to being not tall—but when?

Here are two answers to this question. The supervaluational answer says that there is some time t_i which was the last moment at which Terry was tall, but that for each time t_i, it is not the case that *it* was the last moment at which Terry was tall. The fact that this sounds like a contradiction has motivated an alternative, epistemicist, answer, which agrees that there is some time t_i which was the last moment at which Terry was tall, but holds merely that we do not, and cannot, know which moment that is.

The epistemicist thus differs from the supervaluationist in holding that not only is there a last moment at which Terry was tall, but that someone who says, at each moment, *this is the last moment at which Terry is tall*, would at some point say something true. Let t_n be that moment, and H_{t_n} be Terry's height in inches at that moment. The epistemicist is committed to accepting that:

(1) Someone is tall iff their height is H_{t_n} inches or greater.

Furthermore, there seems to be nothing more to being tall than having a height greater than or equal to H_{t_n}; as Williamson puts it, the vague strongly supervenes on the precise (see Williamson (1994): 202). Given this assumption, not only is (1) true, it is necessarily true—it states the exact conditions under which anyone is, or would be, tall.

So, the epistemicist is committed not just to unknowable contingencies, but also to a striking number of unknowable necessities. Thus, crucial to assessing the outlook of the epistemicist strategy is evaluating her prospects of making sense of these unknowable necessary truths. And, while it is correct that unknown necessities are a fact of life (as Kripke and others have helped us see clearly), it is not enough for the epistemicist to simply claim companions in guilt with our ignorance of mathematical truths (for example), for it is not obvious that the kinds of strategies for explaining unknown necessities in those domains will work equally well to explain our ignorance about the cutoffs of vague predicates.

To see why, notice that a common strategy for capturing ignorance about mathematical truths appeals to impossible worlds where (for instance) certain necessary truths fail to hold. However, such a strategy will struggle to predict that someone ignorant of the precise height cutoff for tallness can believe that tallness facts strongly supervene on precise height facts.

To see why, suppose you know that Jones's height is between H_{t_n} and $H_{t_n} + .01$ inches. Still, plausibly, you don't know that Jones is tall, because his height is too close to the cutoff. Therefore, worlds in your belief state will include:

- w_1: Jones's height is H_{t_n} inches and he is tall.
- w_2: Jones's height is H_{t_n} inches and he is not tall.

Given the necessity of (1), w_2 is an impossible world. But this pair of worlds together violate the principle that tallness facts strongly supervene on height facts, and hence, you don't believe that tallness facts strongly supervene on precise height facts. For an epistemicist who wants to maintain both ignorance about cutoffs as well as believe her theory (which includes the strong supervenience of the vague on the precise), this is an unwelcome result.

As a matter of fact, the most prominent defender of epistemicism, Timothy Williamson, does not explain the unknowability of the cutoff thresholds of vague predicates by appealing to impossible worlds. Rather, he appeals to an extension of the safety condition of knowledge, which he calls *metalinguistic safety*. Whereas the safety condition says that what distinguishes knowledge from merely true belief is that the believer isn't lucky—couldn't easily have been wrong—metalinguistic safety says that knowledge that p requires that the believer be free of metalinguistic luck—that is, that the sentence she uses to express her belief that p couldn't easily have meant something false. Although (1) in fact expresses a necessary truth (we're supposing), we are metalinguistically lucky in expressing a necessary truth with it, since it could have easily expressed a necessary falsehood even though we still endorsed it; this is why we do not know (1).

In this paper, I argue against metalinguistic safety as a necessary condition on knowledge. This puts pressure on epistemicists to find some other way to explain our ignorance of the cutoffs of vague predicates. However, I argue that there is a way out: the epistemicist can deny that the vague strongly supervenes on the precise. I articulate an alternative version of epistemicism that appeals to diagonalized contents that combine metalinguistic and first order material. The primary advantage of diagonalized epistemicism is that it requires only regular (non-metalinguistic) safety to account for our ignorance of cutoffs, and I argue that the cost of denying strong supervenience of the vague on the precise is one worth paying for it. However, in a twist, I argue that even diagonalized epistemicism is false. It is not a matter of luck (epistemic or linguistic) that we do not know claims like (1), since such claims express no facts to be known (or not) at all.

1 Metalinguistic Safety

Williamson's strategy for explaining our ignorance of the cutoffs of vague predicates involves into two independent commitments: (i) a commitment to the semantic plasticity of vague predicates, and (ii) a commitment to metalinguistic safety as a necessary condition of knowledge.

An expression t is semantically plastic iff slight changes in its use lead to slight changes in its meaning. Williamson contends that vague predicates are semantically plastic in this sense:

> For any difference in meaning, there is a difference in use ... A slight shift along one axis of measurement in all our dispositions to use 'thin' would slightly shift the meaning and extension of 'thin'. On the epistemic view, the boundary of 'thin' is sharp but unstable. (Williamson 1994: 231)

Thus, suppose that Jones is exactly n inches tall. There will be a nearby world in which Jones is also exactly n inches tall but in which the usage facts of "tall" differ slightly, leading to the cutoff for "tall" to be slightly different, such that Jones is just barely over the actual cutoff for "tall" but not over the cutoff for "tall" at these nearby worlds. The sentence "Jones is tall" is thus actually true but very easily could have expressed a false proposition. It follows that, if I accept the sentence "Jones is tall," there is a very real sense in which I am linguistically lucky—the words just happened to express a truth for me, but they very well might not have (and I would have continued to accept the sentence).

This kind of linguistic luck is largely uncontroversial. What I *will* challenge as controversial is Williamson's move to hold that my being linguistically lucky with the sentence "Jones is tall" also undermines my knowledge that Jones is tall:

> On the epistemic view, an utterance of a vague sentence such as 'n grains make a heap' may express a necessary truth in a borderline case. A speaker who made such an assertion would not be expressing knowledge that n grains make a heap, for he might easily have used those words even if their overall use had been slightly different, so that they expressed a necessary falsehood. His utterance u does not manifest a disposition to be reliably right. (Williamson 1994: 235)

This suggests the following:

> **Metalinguistic Safety (First Pass)**
> S's utterance of sentence 'P' expresses knowledge that P only if it couldn't have easily been the case that S uttered 'P' and it meant something false.

This principle, together with the semantic plasticity of vague expressions, predicts that when Jones is just barely over the cutoff for "tall," an utterance of a sentence like "Jones is tall" would not express knowledge. Similarly, supposing the cutoff for "tall" is being n inches tall. Then, still, an utterance of the sentence:

(2) Someone is tall iff their height is n inches or greater.

would not express knowledge, since this sentence very easily could have expressed a proposition that was false (say, at a world where the cutoff for "tall" is $n + .01$ inches).

1.1 Refining

An immediate problem with Metalinguistic Safety (First Pass) is that it says only when *utterances* of sentences *express* knowledge, but we are also interested in when some individual knows some proposition, whether or not they utter any sentence that expresses it. Take Smith, who has formed the belief that Jones is tall by measuring him and finding his height to be n inches, yet hasn't uttered any sentences. The epistemicist should still want to predict that Smith doesn't know Jones is tall even though his belief is true and Smith knows that Jones is n inches tall.

Anticipating this kind of thought, Williamson suggests that what goes for sentences goes for concepts too, remarking:

> "The vagueness of an expression consists in the semantic differences between it and other possible expressions that would be indiscriminable by those who understood them. Similarly, the vagueness of a concept consists in the differences between it and other possible concepts that would be indiscriminable by those who grasped them." (Williamson 1994: 237)

> "You have no way of making your use of a concept on a particular occasion perfectly sensitive to your overall pattern of use, for you have no way of surveying that pattern in all its details. Since the content of the concept depends on the overall pattern, you have no way of making your use of a concept on a particular occasion perfectly sensitive to its content. Even if you did know all the details of the pattern (which you could not), you would still be ignorant of the manner in which they determined the content of the concept." (Williamson 1994: 231-2)

In fact, it's not easy to reconcile these two thoughts. The first remark suggests that what's indiscriminable is not the relation of a vague expression to its concept, but rather various distinct concepts themselves. However, the second remark suggests that what's indiscriminable is what content a particular concept (expressed by some vague expression) has. I am not sure which view Williamson ultimately has in mind, but I will opt to flesh out the second view here (I suspect similar remarks will apply to the first).

Here is a way to make Williamson's second suggestion more precise. Suppose that an agent's occurrent beliefs are grounded in mental representations that play some role in their cognitive economy and say that the relation an agent bears to such mental representations is the belief* relation. Suppose also that mental representations have as their meanings propositions (something like this view is endorsed by Fodor 1981, 1987). Thus,

> X believes that p iff there is a mental representation S such that X believes* S and S means that p.

When X believes that p, call the corresponding mental representation S that X believes*, and which means that p, the *mental representative* of X's belief. Then, we can reformulate metalinguistic safety as follows:

> **Metalinguistic Safety (Second Pass)**
> X knows that p only if X truly believes that p and the mental representative of X's belief that p could not have easily meant something false and still been believed* by X.

Supposing that the mental representatives of beliefs involving vague predicates are similarly semantically plastic—such that at nearby worlds, slightly different applications of the mental representatives lead to them having slightly different meanings—and we again predict the result we wanted, without appeal anywhere to utterances of sentences: Smith doesn't know that Jones is tall (even though he truly believes this) and we don't know that someone is tall iff their height is n inches or greater even though we truly believe it.

1.2 Problems

Nonetheless, there is a further problem facing this approach, which stems from the fact that we often have multiple mental representatives underlying the same belief. Take the proposition that Jones is tall. The epistemicist endorsing the supervenience of the vague on the precise must hold that this proposition has the same truth conditions as the proposition that Jones is at least n inches tall (where n is the actual cutoff for tallness). It is no part of the epistemicist's view that we cannot come to know that Jones is at least n inches tall—this is something we can come to know by measuring him for instance. Suppose Smith does come to know this in this way. Then, there is some mental representative underlying this belief, call it S_1, such that:

(3) a. Smith truly believes* S_1, and
b. It couldn't easily have been the case that S_1 meant something false and Smith believes* S_1.

However, in order for Smith not to thereby know that Jones is tall, it must then be the case that there is some distinct mental representative underlying this belief, call it S_2, such that:

(4) a. Smith truly believes* S_2, and
b. It could easily have been the case that S_2 meant something false and Smith believes* S_2.

But the problem is that *what* Smith believes is the same in both cases. Thus, there is no single mental representative underlying Smith's belief, and this causes trouble for Metalinguistic Safety (Second Pass), which presupposes each belief has a unique mental representative. Thus, the view predicts that the claim that Smith knows that p is either undefined or false (depending on their view of uniqueness failures of definite descriptions).

In response, we might modify Metalinguistic Safety (Second Pass) to existentially quantify over mental representatives:

Metalinguistic Safety (Third Pass)
X knows that p only if X truly believes that p and **there is some** mental representative of X's belief that p could not have easily meant something false and still been believed* by X.

However, this yields the result that Smith knows Jones is tall. Alternatively, we might universally quantify over mental representatives:

Metalinguistic Safety (Fourth Pass)
X knows that p only if X truly believes that p and **every** mental representative of X's belief that p could not have easily meant something false and still been believed* by X.

But this yields the result that Smith doesn't know that Jones is at least *n* inches tall.

Another possibility is to relativize knowledge to a mental representative:[1]

Metalinguistic Safety (Fifth Pass)
X knows that p relative to mental representative M only if X truly believes that p and M could not have easily meant something false and still been believed* by X.

It is unclear to me whether knowledge relative to a mental representative bears enough similarities to our ordinary notion of knowledge to warrant serious consideration as an account of the latter. However, let's suppose I'm wrong about this. Nonetheless, it still seems wrong that knowing that p depends on some metalinguistic facts about the representational vehicle M through which we believe that p, such that the semantic plasticity of M could undermine your knowledge of beliefs you have via M.

We can assess Metalinguistic Safety (Fifth Pass) by exploring whether other expressions are semantically plastic in ways not tracked by our dispositions to apply them, and, if so, whether that undermines knowing via beliefs got through mental representatives involving those expressions. For instance, Burge (1979) proposes that the meaning of *arthritis* for a speaker S may depend on the dispositions of his broader social community. Consider the following inversion of Burge's classic case. Given the actual community dispositions, *arthritis* means an inflammation of the joints. Suppose Smith picked up on the use of the term from unreliable sources—suppose he got it from Dr. Jones, someone who always mixes up medical terminology for joints and muscles. In the actual case, Dr. Jones got it right—applying

[1]Even another is to distinguish propositions by mental representatives; I'll set this one aside for now.

arthritis to inflammation of the joints—and so Smith's usage is correct. Smith feels inflammation in his right knee and thus on that basis comes to believe (and know, it seems) that:

(5) I (Smith) have arthritis in my right knee.

But the broader community might have used *arthritis* differently, so that it applied not to inflammation of the joints but rather muscle inflammation. At some such nearby possibilities, when Dr. Jones meets with Smith, he uses *arthritis* correctly (at that world) to talk about muscle inflammation. But at other such nearby possibilities, when Dr. Jones meets with Smith, he uses *arthritis* incorrectly (at that world), applying it to joint inflammation only (remember, he tends to mix up medical terminology for joints and muscles):

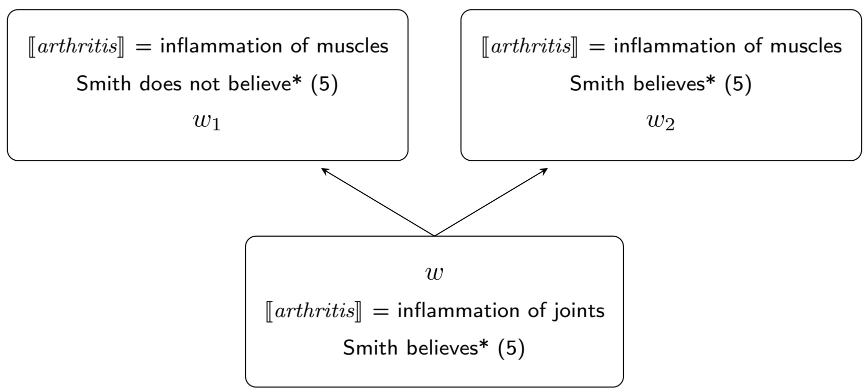

Focus on the latter kind of world, w_2. At that world, Smith's dispositions to use *arthritis* are incorrect (given what it means). As such, since Smith continues to have inflammation in his knee joint at w_2, he continues to believe* the mental representative of his belief that he has arthritis in his right knee. And at w_2, this mental representative expresses a false proposition—the proposition that Smith has muscle inflammation in his right knee. Thus, by Metalinguistic Safety (Fifth Pass), Smith does not know (at w) that he has arthritis in his right knee.

But that result seems wrong. Whether Smith knows that he has arthritis in his right knee doesn't depend on whether the word *arthritis* (or the corresponding mental representative word) could have easily meant something different in a way not tracked by Smith.[2] Knowledge concerns the content of the belief, not the vehicle for it.

This case was designed to mimic a standard fake barn case, which the usual safety condition on knowledge aims to make sense of. In the fake

[2]A similar case is discussed in Bacon (2018): 81.

barn case, Smith is driving through an area with nine fake barns and one real barn. Not recognizing the difference between the fake and real barns, he sets his sights on the one real barn and on that basis comes to believe that there is a barn in the field. Smith does not know there is a barn in the field because he very easily could have believed falsely that there was a barn in the field (by, say, forming the belief by looking at a fake barn). The metalinguistic case substitutes possible alternative meanings for *arthritis* in place of fake barns, but the structure of the cases seems otherwise analogous.

Here is a second example, due to Kearns & Magidor (2008).

> Smith is again in fake barn country and again sets his sights on the one real barn in the vicinity. He walks up to it and inspects it thoroughly from all sides, and concludes, pointing to it, "That is a barn."

I think it is incredibly plausible to think that, in such a case, Smith knows that that thing is a barn. His vision is good, he is paying close attention, has seen it from all angles, and thus he has a true justified de re belief of that thing that it is a barn—had it not been a barn, he would have not believed it to be one.

However, given metalinguistic safety, Smith does not know that that is a barn, since his mental representative could have easily expressed a false proposition. These cases together put a lot of pressure on metalinguistic safety. The lesson I think we should draw is that the semantic properties of the mental representations underlying our beliefs simply do not matter to knowledge: metalinguistic safety is wrong.[3]

2 Supervenience, weak and strong

Let's switch gears to an alternative explanation of the lack of knowledge of borderline cases of vague predicates. If the vague doesn't supervene on the precise, then two possible things could be precisely alike but not vaguely alike, and thus we could know how they are in their precise respects without knowing how they are in their vague respects. Williamson denies this, holding instead that the vague strongly supervenes on the precise:

> "If two possible situations are identical in all precisely specified respects, then they are identical in all vaguely specified respects too. For example, if x and y have exactly the same physical measurements, then x is thin if and only if y is thin." (Williamson 1994: 202)

[3] For additional arguments against metalinguistic safety, see: Mahtani (2004), Caie (2012), Sennet (2012).

We can formulate this principle as follows:

> **Strong Supervenience**
> Necessarily, for any vague property F there is some precise property G such that necessarily, something is F iff it is G.

This says that at every world, every vague property has some precise correlate such that necessarily, the two are co-instantiated. However, by ruling out the possibility of a possibly differing from b in its vague respects but not its precise respects, commitment to strong supervenience forces the epistemicist to explain our ignorance of the vague either metalinguistically or hyperintensionally.

However, notice that we can still capture a sense in which the vague supervenes on the precise without running into this problem. Contrast Strong Supervenience with Weak Supervenience:

> **Weak Supervenience**
> Necessarily, for any vague property F there is some precise property G such that something is F iff it is G.

This says that at every world, every vague property has some precise correlate such that the two are co-instantiated. The difference with Strong Supervenience is that the requirement on co-instantiation is world-bound—thus, which precise property is the supervenience base for some vague property may vary from world to world. Notice, then, that ignorance about which precise G makes something F is ordinary first order factual ignorance, and it could account for our ignorance about the cutoff of vague predicates. This is so, even though at any world, there can be no difference in an object's vague properties without a difference in its precise properties—that is, throughout the class of worlds that agree on the precise supervenience base for some vague property, varying the object's vague properties requires varying their precise properties (in other words: you can never find two actual people who share the same height but differ in whether they are tall).

Just as worlds can vary in their distribution of precise properties, they can differ in their supervenience base of the vague. Here's a graphic to illustrate the difference.

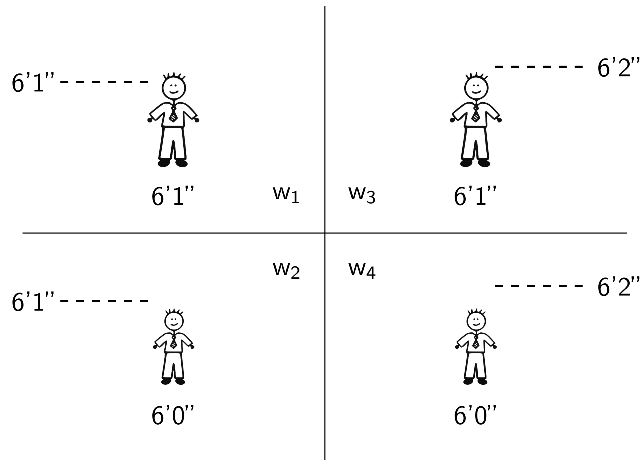

Figure 1

At w_1 and w_2, the cutoff for tallness is 6'1", while at w_3 and w_4 the cutoff for tallness is 6'2". Meanwhile, Jones's height is 6'1" at w_1 and w_3, while his height is 6'0" at w_2 and w_4. These worlds illustrate the dual aspect of our ignorance of vague properties: we are ignorant both about the precise—Jones's precise height—as well as the vague—whether Jones is tall. Furthermore, learning Jones's precise height is not sufficient for learning whether Jones is tall. Learning that, say, Jones is 6'1" would allow us to rule out worlds w_2 and w_4, but we still wouldn't know whether Jones is tall, since we wouldn't know whether being 6'1" is sufficient for being tall.

Finally, if we were to come (somehow) to believe that Jones is 6'1" and that he is tall, we would believe ourselves to be in a world like w_1. But we still wouldn't know that Jones is tall, since our belief isn't safe—there is still a nearby possible world where we believe that Jones is tall, and Jones is 6'1", but the cutoff for being tall is a bit higher, say being 6'2" (a world like w_3 fits the bill).

The extension of a vague predicate like *tall* depends on two factors: the precise height facts, and a minimum cutoff. For the Williamsonian epistemicist, the language fixes the minimum cutoff, and thus the content of *tall* can be specified by a unique precise property of heights. At worlds where the language fixes a different cutoff, the content of *tall* will be a distinct precise property of heights. By contrast, on our alternative version of epistemicism, the content of *tall* incorporates the minimum cutoff, and thus does not vary across worlds with different cutoffs.[4]

Looking back to Figure 1 above, suppose for a moment that w_1 is the actual world. Then, according to the Williamsonian epistemicist,

[4] I count Barker (2002, 2013) among epistemicists who endorse this latter kind of epistemicism, although I'm not sure Barker would agree with my interpretation of his view. See also MacFarlane (2020) for a characterization of diagonalized epistemicism that matches the version I describe here.

(6) Jones is tall.

(7) Anyone at least 6'1" is tall.

(8) It is necessarily true that anyone at least 6'1" is tall.

By contrast, according to our alternative view, (6) and (7) are true, but (8) is false: it is contingent fact that anyone at least 6'1" is tall.

3 Diagonalizing

So far, we've followed Williamson in assuming, plausibly, that it is not essential to a language that its meaningful parts have the meanings they do have—they might have meant different things had the facts relevant to their metasemantics been different. Thus, we can define a possibly partial interpretation function $[\![\ ]\!]$ that assigns a propositional content to a sentence relative to a possible world. So,

$$[\![A]\!]_w = \text{the content of } A \text{ at } w.$$

Given our assumption that propositional contents can be modeled as sets of possible worlds, we have that:

$$[\![A]\!]_w^{w'} = 1 \text{ (true) iff the content of } A \text{ at } w \text{ is true at } w'.$$

For terminology, we'll say that the world that determines the expression's content is the **determining** world, while the world we evaluate the resulting content relative to is the **evaluation** world. So, above, w is the determining world and w' the evaluation world. Next, we'll suppose for illustration that vague predicates like *tall* have a covert variable whose value is fixed by the determining world:

$$[\![\textit{Jones is tall}_s]\!]_w^{w'} = 1 \text{ iff Jones's height is at least } [\![s]\!]_w \text{ at } w'.$$

With this notation in hand, we can state Williamsonian epistemicism and our alternative. Let α be the actual world. Then:

Williamsonian Epistemicism
The content of *Jones is tall* at α is $\{w\colon [\![\textit{Jones is tall}_s]\!]_\alpha^w = 1\}$

In other words, the set of worlds where Jones's height meets the actual standards of tallness.

Diagonalized Epistemicism
The content of *Jones is tall* at α is $\{w\colon [\![\textit{Jones is tall}_s]\!]_w^w = 1\}$

In other words, the set of worlds w where Jones's height at w meets the standards of tallness at w.

What looks like an incredibly small difference here in fact makes *all* the difference for avoiding the challenges put forward in §1. To see this, let's first see why these views agree about (6) and (7) but disagree about (8).

Recall: Jones is actually 6'1" and 6'1" is the minimum height to count as tall. Thus, both views predict *Jones is tall* is true and that *Anyone who is at least 6'1" is tall* is true. This is unsurprising, since both views hold that *Jones is tall* is true at α iff Jones's height at α meets the standards of tallness at α.

Things change when the content of *Anyone who is at least 6'1" is tall* is embedded. According to Williamsonian epistemicism, the determining world plays no role after fixing the content of *Anyone who is at least 6'1" is tall*. Once fixed by the actual world, the content is $\{w$: $[\![$*Anyone who is at least 6'1" is tall*$_s]\!]^w_\alpha = 1\}$. Then, when embedded under a modal operator like "Necessarily" we have:

$[\![$*Necessarily, anyone who is at least 6'1" is tall*$_s]\!]^\alpha_\alpha = 1$ iff
at all worlds w : $[\![$*Anyone who is at least 6'1" is tall*$_s]\!]^w_\alpha = 1$ iff
at all worlds w : Anyone whose height is at least 6'1" at w has a height of at least $[\![s]\!]_\alpha$ at w iff
at all worlds w : Anyone whose height is at least 6'1" at w has a height of at least 6'1" at w

And thus we predict that (8) is true.

(8) It is necessarily true that anyone at least 6'1" is tall.

By contrast, according to diagonalized epistemicism, the determining world shifts alongside the evaluation world when the content of *Anyone who is at least 6'1" is tall*. Thus,

$[\![$*Necessarily, anyone who is at least 6'1" is tall*$_s]\!]^\alpha_\alpha = 1$ iff
at all worlds w : $[\![$*Anyone who is at least 6'1" is tall*$_s]\!]^w_w = 1$ iff
at all worlds w : Anyone whose height is at least 6'1" at w' has a height of at least $[\![s]\!]_w$ at w'

The derivation ends here, since $[\![s]\!]_w$ depends on the world w that has been shifted by the modal operator. Since this condition won't hold for all worlds, we predict that (8) is false.

We can also see now why Williamsonian epistemicism entails Strong

Supervenience, whereas diagonalized epistemicism only entails Weak Supervenience:

> **Strong Supervenience**
> Necessarily, for any vague property F there is some precise property G such that necessarily, something is F iff it is G.
>
> **Weak Supervenience**
> Necessarily, for any vague property F there is some precise property G such that something is F iff it is G.

Our candidate vague property is tallness. Recall the model from above:

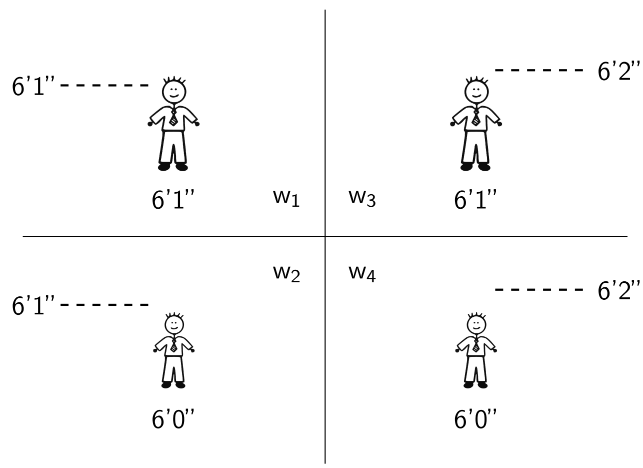

Figure 1

Let's suppose w_1 is the world of utterance. Then, we can calculate the truth values of the contents of *Jones is tall* at each of these worlds, given Williamson epistemicism and diagonalized epistemicism. Recall the difference in their predicted contents:

> Williamson Epistemicism: $\{w : [\![\textit{Jones is tall}_s]\!]^{w}_{w_1} = 1\}$
>
> Diagonalized Epistemicism: $\{w : [\![\textit{Jones is tall}_s]\!]^{w}_{w} = 1\}$

For the Williamsonian Epistemicist, the precise property necessarily co-extensive with *tall*$_s$ as uttered in w_1 is being at least 6'1" tall. And so, two possible individuals (Jones at w_1 and Jones at w_2, for instance) can differ in whether they are tall only if they differ in their heights. By contrast, for the diagonalized Epistemicist, there is no precise property that is necessarily co-extensive with *tall*$_s$ uttered in w_1. And, thus, there can be two possible individuals (Jones at w_1 and Jones at w_3) who differ in whether they are tall but do not differ in their heights; thus, the view predicts violations of Strong Supervenience. However, the diagonalized Epistemicist still predicts that at any world, no two individuals can differ in whether they are tall without differing in their height at that world—this is because the cutoff for tallness is a property of the world. Thus, the view predicts Weak Supervenience.

4 Motivating Diagonalized Epistemicism

Suppose you are already a committed epistemicist—that is, you think that vague predicates have precise but unknowable cutoffs—but you don't want to endorse metalinguistic safety as the explanation for our ignorance of such cutoffs. Then, you have a reason to be a Diagonalized Epistemicist. Your ignorance of the cutoffs of vague predicates is just ordinary factual ignorance that can be predicted by the standard safety condition on knowledge.

Here's another argument in favor of diagonalized epistemicism over Williamsonian epistemicism. The argument draws on Stalnaker's theory of communication (see Stalnaker 1978, 1999, 2002). For Stalnaker, to assert a proposition p is to propose making p jointly believed (or taken for granted, for the purposes of the conversation). We do this by uttering a sentence S with the requisite assertoric intentions in the right context (where we can reasonably believe to be taken seriously). If all goes well, that sentence expresses a unique proposition p, which is recognized by our interlocutors to be the proposition we are proposing to be jointly accepted. If we are in fact sincere and they trust us, then the proposition will become jointly accepted between us.

However, when an utterance of S occurs in a context in which there is joint ignorance about what S's propositional content is, it will be unclear to hearers just what belief they should adopt. Suppose Jones utters S to Smith, but Smith doesn't know whether S means p or S means q. Then, Smith won't know whether Jones has proposed that they jointly believe p or jointly believe q, and hence Jones's assertion will be infelicitous.[5] This is borne out in some cases:

(9) a. *Doctor to patient*: You are suffering from acute myositis of the extensor retinaculum and require an arhroscopic myectomy to fully recover.
    b. *Patient*: What??

Williamsonian epistemicism predicts, for this reason, that some assertions

[5]This reasoning motivates Stalnaker's Uniformity constraint (Stalnaker 1978: 325):

> **Uniformity**
> Utterance U expresses the same proposition at every world compatible with what is mutually presupposed in the context.

of vague sentences should warrant such a response.[6] Consider the following context:

> Alice is a fifth grade basketball coach, and her friend Beth is a college basketball coach. Beth never spends time around fifth graders, so she has no idea how tall they are on average. The two are discussing a student, Cathy, and whether she would be a good fit on Alice's team of fifth graders. Beth knows Cathy's on-court stats, but not her height or whether she is taller on average than other fifth graders, whereas Alice knows the latter but not the former. Finally, let's suppose they know these facts about each other.

For simplicity, we'll suppose that for all Beth knows, every combination of the following are possible:

- For each $n : 4'0'' \leq n \leq 7'0''$: Cathy is n feet tall.
- For each $i : 5'0'' \leq i \leq 6'5''$: A fifth grade basketball player must be at least i feet tall to be tall.

Now, suppose Alice tells Beth:

(10) Cathy is tall.

Williamsonian epistemicism predicts that Beth should regard Alice's claim with the same confusion as the patient responding to the doctor above. Beth has no idea what proposition Alice has expressed, and thus no idea what proposition has been proposed for her to believe.[7]

However, that is not how Beth should respond to Alice's claim. Remember, Beth knows that Alice is better positioned than she is to know what the standards of being tall for a fifth grade basketball player are. Thus, it seems much more plausible that Beth would respond by accepting Alice's assertion and thus rule out possibilities in which Cathy's height is below the minimum standards for what it takes to be tall for a fifth grade basketball player. In other words, she rules out any possibility in which:

> Cathy is n feet tall and a fifth grade basketball player must be at least i feet tall to be tall and $n < i$.

[6] Of course, they *need not* warrant such a response. For instance, when the possible propositional contents p and q are equivalent given what's commonly believed, the hearer can recognize without harm to the speaker's intentions that the belief proposed is $p \cup q$. The cases we are considering here are not of this kind.

[7] This case presents a different challenge to Williamsonian epistemicism than the arthritis case from §1.2. In the arthritis case, the agent doesn't know that they have arthritis in their knee because the mental representative of that belief could have meant something (they know) it does not. In this case, Beth doesn't know what claim Alice has made, and thus doesn't know how to update her beliefs.

This is not predicted by Williamsonian epistemicism, but it is predicted by Diagonalized epistemicism (see also Barker 2002, MacFarlane 2020 for similar arguments). Count one more for Diagonalized epistemicism.

5 Against Epistemicism

I've argued so far that, if you are an epistemicist about vagueness, you should be a diagonalized epistemicist. But, alas, now we come to the rub: I think you shouldn't be an epistemicist about vagueness. Let's start with a case that might, at first sight, be taken to be friendly to the epistemicist (adapted from Schiffer 2000: 223-4).[8] Remember Terry, who is 6'7" and is clearly tall, and who is about to be zapped by the Shrink Ray 3000 (and thus be shrunk by a quarter inch per second for one minute). Since Terry is clearly tall at t_0, we can suppose the following:

> At t_0,
> i. You are certain that Terry is tall.
> ii. It would be correct to assert that Terry is tall.

After 28 seconds in the shrink ray (at t_{28}), when Terry is now 6'0" and, we may suppose, clearly a borderline case of being tall, it seems that:

> At t_{28},
> i. You are not certain that Terry is tall.
> ii. It would not be correct to assert that Terry is tall.

Between t_0 and t_{28}, what happens? According to Schiffer, you slowly lose confidence that Terry is tall, and, if asked whether Terry is tall, should respond with qualifications ("Well, he kind of is."). Nicholas Smith concurs, noting that, as Terry shrinks, you would become "less and less sure" that Terry is tall, and "more and more sure" that he isn't (see Smith 2010: 3).

Are these claims correct? Is it correct to be less and less sure that Terry is tall as he shrinks? Diagonalized epistemicism says yes: as Terry shrinks, since you aren't sure what the standards for tallness are, whether his height meets the standards of tallness becomes less and less likely. However, I think your subjective probability that Terry is tall does not decrease as he shrinks. Rather, I will argue that, as he shrinks, it eventually ceases to be the case that you ought to be certain that Terry is tall. But this is not because you ought to be less than certain that Terry is tall; rather, it is because there is no $n : 0 < n < 1$ such that it is determinate that you think it is n-likely that Terry is tall.

[8] See also Smith (2009, 2010).

To see why, consider a variant case. Suppose you are told to estimate the probability that the ball drawn from (and then put back into) a jar was striped. At first, you're told by a reliable source that the ball was selected from Jar 0, which you can see contains only striped balls. In that case, the following seem correct:

At J_0,

i. You are certain that the ball was striped.

ii. It would be correct to assert that the ball was striped.

But now suppose you're told that, actually it was chosen from (and put back into) Jar 1, which you can see contains all striped balls except for one. In that case, it seems that:

At J_1,

i. You are not certain that the ball was striped.

ii. It would not be correct to assert that the ball was striped.

The reason it would not be correct to assert that the ball was striped is that there is a salient possibility that you can't rule out—namely that the ball drawn was the one non-striped ball.

Suppose this process is iterated so that at each stage you're told, actually, it was chosen from Jar n, where you can see that the proportion of striped to non-striped balls is lower than the previous stage (we might imagine your memory is wiped so you may continue to believe your information source is reliable). Here, I think it is clear that over the course of this operation, your subjective probability that the ball was striped should decrease overall (whether it will decrease monotonically will depend on your eyesight and ability to estimate proportions of striped/non-striped balls), and it would remain incorrect to assert that the ball was striped. By the point at which there is a close to even ratio of striped to non-striped balls in the jar, you should think it around .5 likely that the ball was striped. These judgments are quite clear.

Compare these probability judgments with those in Terry's case. At t_{28}, Terry is 6'0"—a clear borderline case of being tall (we've supposed). We've agreed that you shouldn't be certain that he's tall, nor certain that he's not tall. Suppose we pause the shrink-ray at this point and I ask you how likely it is that Terry is tall. Here, while the (A) response sounds odd, the (B) and (C) responses are perfectly acceptable:

(A) #He's (about) as likely tall as not.

(B) What do you mean? He's only kind of tall.

(C) Well, I'm not certain he's tall.

The infelicity of (A), especially when compared to the acceptability of (B) and (C), suggests that you should not think Terry is around .5 likely tall. However, the acceptability of (C) confirms the initial Schiffer/Smith intuition that:[9]

> **Indeterminacy undermines certainty**
> If you are certain that it is indeterminate whether S is F, then it is determinate that you are not certain that S is F, nor certain that S is not F.

But the combination of (A) and (B) reveals something further. Given that it is indeterminate that S is F iff S is only kind of F, we can motivate the following principle:[10]

> **Indeterminacy undermines (determinate) probability**
> If you are certain that it is indeterminate whether S is F, then for any $n : 0 < n < 1$, it is indeterminate whether you think S is n-likely F.

One argument for this conclusion is that strings of the following form are infelicitous:

(11) S is only kind of F, and in fact S is n-likely F.

a. #John is only kind of a jerk, and in fact he's $\left\{\begin{array}{l}\text{likely}\\ \text{not likely}\\ \text{as likely as not}\end{array}\right\}$ a jerk.

b. #That table only kind of flat, and in fact it's $\left\{\begin{array}{l}\text{likely}\\ \text{not likely}\\ \text{as likely as not}\end{array}\right\}$ flat.

This infelicity mirrors the infelicity exhibited by:

(12) S is only kind of F, and in fact S is / is not F.

[9] This principle, and the one following, provide constitutive connections between credences in indeterminacy with determinate/indeterminate credences. An alternative conclusion to draw is that what's motivated here are principles about what credence you ought to have, given certain credences in indeterminacies. I will set aside the distinction between these types of principles here.

[10] The shift to the "kind of" locution is to avoid complications arising from the fact that "indeterminate" isn't obviously a term of ordinary English, and hence one that we can rely on pre-theoretical intuitions about.

a. #John is only kind of a jerk, and in fact he is / isn't a jerk.
b. #That table only kind of flat, and in fact it is / isn't flat.
c. #Terry is only kind of tall, and in fact he is / isn't tall.

Just as the sentences in (12) are infelicitous because one ought not assert that p if one is sure that it is indeterminate whether p, so it is that one ought not assert that it is n-likely that p when one is sure that it is indeterminate whether it is n-likely that p. Thus, given Indeterminacy Undermines (Determinate) Probability, we explain why the sentences in (11) are infelicitous too. And thus, contra Schiffer and Smith, as Terry shrinks, it would not be correct to be less and less sure that he is tall. Nonetheless, this view can explain what might have motivated the Schiffer/Smith intuition: given Indeterminacy Undermines Certainty, you are not certain that Terry is tall when you're sure he's only kind of tall.

Unfortunately, epistemicism does not predict Indeterminacy Undermines (Determinate) Probability. Instead, the diagonalized epistemicist predicts that when you are certain that it is indeterminate whether S is F, there will be some $n : 0 < n < 1$ such that, determinately, you think S is n-likely F.[11] Take Terry's case at t_{28}, where you know he is 6'0" but (let's say) are uncertain whether being 6'0" is sufficient to be tall. Idealizing, let's suppose you're uncertain between two possible cutoffs for tall: 6'0" and 6'1". Then, the diagonal of *Terry is tall* is:

$$\{w : [\![\textit{Terry is tall}]\!]_w^w = 1\}$$

and this proposition will be true at w_1, where the cutoff is 6'0" and Terry is 6'0", and false at w_2, where the cutoff is 6'1" and Terry is 6'0". Since these are the only epistemically possible worlds (again, we are idealizing, so these will be classes of worlds), your probability in the diagonal of *Terry is tall* will be .5, and there is nothing indeterminate about this—contra Indeterminacy Undermines (Determinate) Probability.

That is why I am not an epistemicist.[12]

6 Conclusion

I started with an objection to epistemicism: according to the view, why is it that we cannot know the precise cutoffs of vague predicates? I argued against the standard, Williamsonian, answer, that appeals to metalinguistic

[11] If they think credences are mushy, then n here will be a range.

[12] I leave aside the question how best to capture these judgments. I am optimistic that some linguistic theory of vagueness will do better in this regard, since linguistic theories are well positioned to predict that embeddings of indeterminate sentences under modal operators will yield indeterminate sentences—see for instance, Dorr (2003), Rayo (2008), Sud (2018).

safety. In its place, I suggested that epistemicists should be diagonalized epistemicists. But, then I argued against diagonalized epistemicism. Thus, ultimately, I think the best version of epistemicism fails—we should not be epistemicists about the vague.

References

Bacon, Andrew. 2018. *Vagueness and thought*. Oxford: Oxford University Press.

Barker, Chris. 2002. The Dynamics of Vagueness. *Linguistics and Philosophy*, **25**, 1–36.

Barker, Chris. 2013. Negotiating Taste. *Inquiry*, **56**, 240–257.

Burge, Tyler. 1979. Individualism and the Mental. *Midwest studies in philosophy*, **4**, 73–121.

Caie, Michael. 2012. Vagueness and semantic indiscriminability. *Philosophical studies*, **160**(3), 365–377.

Dorr, Cian. 2003. Vagueness without ignorance. *Philosophical Perspectives*, **17**(1), 83–113.

Fodor, Jerry. 1981. Some notes on what linguistics is about. *Readings in the Philosophy of Psychology*, **2**, 197–207.

Fodor, Jerry. 1987. *Psychosemantics: The problem of meaning in the philosophy of mind*. Cambridge, MA: MIT press.

Kearns, Stephen, & Magidor, Ofra. 2008. Epistemicism about vagueness and meta-linguistic safety. *Philosophical Perspectives*, **22**, 277–304.

MacFarlane, John. 2020. Indeterminacy as Indecision. *The Journal of Philosophy*.

Mahtani, Anna. 2004. The instability of vague terms. *The Philosophical Quarterly*, **54**(217), 570–576.

Rayo, Agustín. 2008. Vague representation. *Mind*, **117**(466), 329–373.

Schiffer, Stephen. 2000. Vagueness and Partial Belief. *Philosophical Issues*, **10**, 220–57.

Sennet, Adam. 2012. Semantic plasticity and epistemicism. *Philosophical studies*, **161**(2), 273–285.

Smith, Nicholas J. J. 2009. *Vagueness and Degrees of Truth*. Oxford: Oxford University Press.

Smith, Nicholas JJ. 2010. Degree of belief is expected truth value. *Pages 491–506 of:* Dietz, R., & Moruzzi, S. (eds), *Cuts and clouds: Vagueness, its nature, and its logic*. Oxford: Oxford University Press Oxford.

Stalnaker, Robert. 1978. Assertion. *Pages 315–332 of:* Cole, P. (ed), *Syntax and Semantics 9: Pragmatics*. New York: Academic Press.

Stalnaker, Robert. 1999. *Context and Content*. Oxford: Oxford University Press.

Stalnaker, Robert. 2002. Common Ground. *Linguistics and Philosophy*, **25**, 701–721.

Sud, Rohan. 2018. Plurivaluationism, supersententialism and the problem of the many languages. *Synthese*, 1–27.

Williamson, Timothy. 1994. *Vagueness*. New York: Routledge.