Superfine Sorites and Epistemicism

Siddharth Muthukrishnan
Harvard University
siddharth_muthukrishnan@fas.harvard.edu

January 26, 2026

Abstract

Epistemicism about vagueness is the view that vague terms, like "bald", have sharp boundaries but we are ignorant of where these boundaries lie. According to Williamsonian epistemicists, these sharp boundaries arise from patterns of use. I develop a novel objection to Williamsonian epistemicism using what I call *superfine sorites*: for any vague term we can construct sorites sequences using arbitrarily fine-grained descriptions—hairs, molecules, quantum mechanical wavefunctions, even states described only by hitherto unknown fundamental physics. The epistemicist must commit to sharp boundaries at all these grains. I argue that epistemicists' use-based metasemantics is too weak to underwrite these sharp boundaries because use patterns do not carry enough detail to fix arbitrarily fine-grained cutoffs. There is a way out for the epistemicist: Take the microphysical details of use-instances to be relevant to a successful metasemantic account. However, taking use to be so fine-grained is ill-motivated from the perspective of a metasemantic theory.

1 Introduction

Vague terms—such as "tall", "bald", and "heap"—are those which do not seem to allow for sharp boundaries to be drawn between when they apply and when they

don't (Sainsbury 1996). This leads to the sorites paradox—e.g., by increasing the number of grains of sand in a collection one by one we can turn a non-heap into a heap even though it seems as if no single grain could've made the difference.

Epistemicism is an approach to the problems of vagueness that says it only *seems as if* vague terms don't have sharp boundaries.[1] Epistemicists say that vague terms *do* have sharp boundaries, but we think they don't because we don't know, and indeed can't know, where those boundaries lie.

While many objections have been raised against epistemicism[2], none have focused on the relation between epistemicism and the structure of our physical world, in particular the fact that the physical world has incredibly detailed structure at scales unimaginably smaller than typical human concerns. In this paper, I show that epistemicism is committed to sharp cutoffs at all those extremely small scales, i.e., even in what I call *superfine sorites* sequences (Sec. 3). I then use that observation to generate a novel objection to epistemicism: their metasemantics is too weak to underwrite such sharp cutoffs (Sec. 4). Finally, I consider potential responses and argue that they are not convincing (Sec. 5).

2 Williamsonian Epistemicism

The version of epistemicism I focus on is Tim Williamson's, set out in his book *Vagueness* (Williamson 1994). The main element of this view we will focus on is the way in which sharp cutoffs are generated.

Williamson appeals to semantic externalism. The meanings of our terms are at least partially grounded in their use. The meaning of a given term that these use patterns influence can sometimes be stabilized by the natural structure of

1. The most prominent epistemicist is Williamson (1994). I will focus on his work. See Sorensen (2001) for a different epistemicist picture.
2. See, e.g., Wright (1995), Schiffer (1999), Keefe (1995), Graff (2002), and Sennet (2012).

the world (Bird and Tobin 2024). Williamson argues that vague terms don't enjoy such stabilization by the world. This is why drawing a boundary for them *seems* arbitrary. Nevertheless, the patterns of use fix a sharp boundary.

The complexity of the way in which patterns of use fix the sharp boundaries helps explain, in part, why we are unable to come to know where these boundaries lie. Thus, on this view, what fixes the precise number n of grains of sand that demarcates the boundary between heaps and non-heaps is the diverse and complex ways in which ordinary speakers have used the term "heap" up to this time.

Another reason why we can't come to know where the sharp boundary is is that even if we did come to have a true belief about where this boundary lies, such a belief would not be safe—i.e., there would be very nearby possible worlds where we could come to have the same belief but it would be false. Why? Because vague terms are *plastic*, i.e., even small differences in the way in which these terms are used can create changes in their meaning, and hence change where the cutoff lies.[3]

3 Superfine Sorites

According to epistemicism, there is a precise number of hairs at which a person turns from being bald to not bald. This blocks the sorites paradox as stated in terms of number of hairs. But there's no requirement that sorites sequences must be constructed only in terms of number of hairs. One could, instead of removing one hair at a time, remove a millimeter of hair at a time, and at some point after removing some length of hair we'd have removed one hair, and then we'd move on to the next hair, and so on, eventually rendering the person bald. This too would generate a sorites sequence. And the epistemicist would have to

3. See Magidor (2019).

say that there was a sharp cutoff on this fine-graining as well.

We could go further: we could remove just one molecule of hair at a time. As we remove one molecule at a time, we'd again transition, according to the epistemicist, at some precise point from a non-bald person to a bald person.

We could set up the iteration at an even finer grain. Molecules are not discrete balls connected with sticks (as one might be led to believe from middle school chemistry). Rather they are quantum mechanical systems, with their constituent electrons and nucleons in some state that is represented by a wavefunction. One could then change the quantum-mechanical wavefunction continuously and eventually transition a non-bald person into a bald person. The epistemicist is committed to a precise transition point on this continuum of wavefunctions.

You get the gist. One can keep adverting to deeper and deeper physical theories: quantum field-theoretic states, and then deeper still, string-theoretic states or loop-quantum-gravity states or the states of whatever theory it is that is more fundamental than quantum field theory. For all we know, there maybe no final theory but just an infinite hierarchy of more fundamental theories—tortoises all the way down. And the epistemicist needs to be committed to there being sharp cutoffs for the non-bald to bald transition in the sorites sequences described at all these (potentially infinite) levels of description.

This, then, is what I'm calling *superfine sorites*: we can set up sorites sequences at arbitrarily fine grains of description, relying on the vocabulary and structures of arbitrarily fine grains of physical description.

4 Why superfine sorites is a problem for the epistemicist

Superfine sorites is a problem for the epistemicist. Not because it makes a counterintuitive view even more counterintuitive; I largely agree with Williamson (2004) that intuitions have little evidential value here. And not because the epistemicist cutoffs in deep physics are not salient to physicists; according to the epistemicist, the source of the sharp cutoffs was always in those aspects of our linguistic practices *not* constrained by the world and hence precisely what would not be salient to physicists. No, superfine sorites is a problem for the epistemicist because the superfine sharp cutoffs must be generated via appeal to patterns of use of ordinary language and it's difficult to see how that can be so.

Williamson doesn't give us a fully articulated metasemantic model of how vague terms get their sharp boundaries. But the core idea is that for vague terms (which are not stabilized by the world) use fixes or determines meaning. Consequently, meaning supervenes on use. Hence, any change in the sharp cutoff of a term must result from some variation or change in patterns of use.

What kind of variations in patterns of use might possibly impact changes in meaning? To get a handle on this, suppose the epistemicism is correct. Now, consider two worlds: one world is the actual world and the other world is very similar to actuality except that the number of hairs that marks the boundary between bald and not bald is slightly larger in that world compared to the boundary in our world.

Because meaning supervenes on use, there must be some difference in patterns of use between the actual world and the other world. Perhaps English speakers in that world tend to be just a tiny bit more permissive in who they call bald. Perhaps the individuals we might classify as *nearly almost bald*, the denizens of

the other world tend to classify as simply *bald*.

But now, with the idea of superfine sorites in hand, consider two worlds that differ with respect to the epistemicist cutoff for *bald*, but not at the level of number of hairs, but at the level of the quantum wavefunction's hundredth decimal place. So this is an extremely fine-grained difference in the extension of *bald* between these two worlds. As before, according to the epistemicist's preferred metasemantics, this difference must be grounded in *some* difference in use.

And here is the core challenge: There is nothing that we could reasonably call *use* that could plausibly make such a fine-grained difference. Patterns of use typically consist of things like utterances, inscriptions on paper or computers, nods or head-shakes, gestures, dispositions to respond, and so on. These just aren't sorts of things that are sufficiently richly detailed to be the reason for a difference that is so incredibly small.

To better see that patterns of use aren't so richly detailed, consider the sorts of things that could drive differences in use. Let's go back to the case where the cutoff for bald in terms of number of hairs differs between two worlds. One might imagine that perhaps the reason why the denizens of the other world are slightly more permissive in who they call "bald" is that they have just a tiny bit weaker ability to discriminate the number of hairs on their head (perhaps the typical artificial lighting in their world is just a tiny bit dimmer). So these kinds of differences in perceptual capabilities or conditions ground differences in patterns of use.

However, if the difference between the cutoffs for bald is only at the extreme microphysical level, then differences in perceptual capabilities or conditions just cannot ground usage differences because the difference is (by hypothesis) far too small to be even remotely salient to human perceptual capabilities or

conditions—aided or unaided.

Here's another way to see that use just isn't that fine-grained. Think of the attribution of a sharp cutoff to a given pattern of use as a kind of *semantic theory* that fits that pattern of use. This is broadly in line with the Quine-Davidson-Lewis metasemantic tradition that the epistemicist is drawing on.[4] In other words, the cutoff should explain something about the patterns of use.

But then that means that superfine variations of cutoffs simply cannot be explanatory of use differences. That is, if you imagine changing the cutoff in say a quantum-field-theoretic state at the hundredth decimal place, this difference must be able to explain some change in use. But there just isn't anything that we could meaningfully call use that such a small change in cutoff could influence.

In sum, superfine sorites is a problem for the epistemicist because the epistemicist wants the superfine cutoffs to be determined by use, but use is too coarse-grained for that task. A signature of this is that superfine variations in meaning cannot be adequately explained by coarse-grained variations in use.

5 Potential responses

One potential response by the epistemicist to the objection from superfine sorites is to argue that not all fine-grained cutoffs are allowed. If the problem was that there just isn't enough ways for use to vary so as to be the supervenience base for all the different cutoffs, then perhaps most cutoffs can't be realized. That is, a given coarse-grained pattern of use will determine an extremely fine-grained cutoff, but that most other cutoffs, including very nearby ones, are not possible because there is no pattern of use that can ground all those different cutoffs. In other words, because meaning supervenes on use, and use can only vary in a

4. The closest we get to a metasemantics from Williamson is in Chapter 8 of Knowledge and its Limits, and that is broadly along this line.

coarse-grained fashion, there just aren't enough ways to vary use so as to make all the different fine-grained cutoffs possible, and consequently, there just can't be superfine variations in meaning.

Note that this line of argument does not really supply an explanation of *how* it is that we are able to generate superfine cutoffs in the first place. That is, it doesn't say how coarse instances of use can ground superfine meaning. It simply defuses the problem that superfine meaning can't plausibly supervene on coarse use.

But as a result, this response faces an explanatory burden: Why are only an extremely small (possibly measure zero) set of cutoffs permitted? And why *these* cutoffs and not ones infinitesimally close to them? Why aren't so many cutoff values that one would have every reason to think are possible values of cutoff ruled out?

I don't mean to suggest that these answers are unanswerable. Instead, having to provide adequate answers to them would be a cost of this view.

A related response the epistemicist could try: A given pattern of use fixes not a specific cutoff but an entire range of cutoffs, any of which is compatible with the pattern of use. But this view would be unattractive for the Williamsonian epistemicist. One natural interpretation of the range-of-cutoffs view is that we are free to pick any cutoff within the range. If so, then whether or not someone counts as bald isn't fixed simply by use but also by an arbitrary choice of a cutoff within the range, and there will be no fact of the matter which is the right choice. This makes epistemicism considerably less attractive.

Another issue with the range-of-cutoffs picture that the Williamsonian epistemicist faces is that it invites the problems of higher-order vagueness into the epistemicists' system, for then they'd have to explain what explains the location of the sharp boundary of this range. In particular, they'd have to explain why

this particular cutoff and not a very nearby one. And this just reiterates the dialectic above. Given that avoiding the problems of higher-order vagueness is a central attraction of epistemicism, a response of this nature would demand too high a price.

In sum, I don't think the epistemicists would try to rule out a whole range of prima facie plausible values of the cutoffs or try to argue that use fixes only a *range* of cutoffs.

A different and more plausible response that the epistemicist might put forward is the following: Why should we take usage to be this coarse-grained thing anyway? Instances of use—utterances, inscriptions, gestures, and so on—are all physically realized and hence one describe them using arbitrarily deep physics; e.g., you can take an utterance and describe it in terms of molecules or quantum fields or strings or whatever. So use, too, has superfine structure.

That instances of use have superfine structure is a straightforward deliverance of our modern scientific understanding of the world. But this deliverance can be employed by the epistemicist in the following way: when patterns of use fix a sharp cutoff, the superfine structure of the use plays a role. This then renders plausible that aspects of meaning—in particular, the cutoff—can have superfine structure. That is, the sharp cutoff for a vague term like "bald" is being determined by use patterns which are themselves exceptionally fine-grained.

And if one takes use-instances to be very fine-grained, then there could be extremely small changes in use—changes happening far below the threshold of human sensation—that would then allow for extremely small variations in meanings. More generally, taking use to be extremely fine-grained allows use patterns to determine extremely sharp cutoffs in meaning.

So this response seems to address the worries raised by the presence of superfine sorites sequences for the epistemicist.

The problem, however, with this view is that it's unclear why a metasemantic theory, the one that connects meaning to use, has to respect the fine-structure of use. There is no motivation, *qua* what constitutes a successful metasemantic account of how a language gets its semantics, to care about the physical fine-structure of use-instances.

When constructing a metasemantic theory, we have reason to look at patterns of assent or dissent, to look at patterns of intonation of voices, to look at how people gesture and tilt their heads, and so on. We have some idea how a succesful metasemantic account would incorporate these details. But there is no reason why we would ever care about the 100th decimal place of some quantity described by some heretofore unknown physical theory; there's no plausible reason why a successful metasemantic account would be sensitive to those, *except* simply to preserve epistemicism. That is, outside of a desire to preserve epistemicism, there's no metasemantic task to which we can put the superfine structure of use instances—there are no interesting metasemantic explanations it features in.

To make this point a bit more salient: When we think of metasemantics—especially when we are trying to understand how the meanings of terms that are *not* stabilized by natural structure—we think of it as a *special science*, that it is not a topic that requires the full-bore vocabulary of advanced fundamental, or close-to-fundamental, physical theories.

More specifically, we think of the core items that metasemantics studies to be *functionally realized* (again insofar as it pertains non-natural terminology). Consider *use*. Suppose in one instance I use the term "heap" by writing the word in a sentence on paper; and in another instance I write the sentence in an email on my computer; and in yet another instance I speak the sentence out loud. As far as being inputs into a metasemantic theory, these different uses shouldn't count as fundamentally different. What matters are the following sort of things:

how the word functions and is related to the other words in the sentence, what sentences appear before or after this one, what speakers and listeners do after uptake of this sentence, and so on. However, if we assume that the superfine structure of use-instances is metasemantically relevant, then these three different kinds of use must count as wildly different because they differ wildly in their microphysical instantiation. This is implausible.

One could respond to this charge of implausibility by saying that superfine structure of use-instances are *not* relevant to the *coarse-grained* semantics of vague terms—i.e., those aspects of meaning relevant in everyday life. But the superfine structure of use-instances *is* relevant to the superfine structure of meaning, which is something that is posited by epistemicism.

But this then brings us back to the point I made above: it isn't clear that there is much work being done by saying that the superfine structure of use is metasemantically relevant *other* than simply rescuing epistemicism. An epistemicist might be willing to accept this, but it means their use-based metasemantics is not independently motivated—it needs to be significantly modified so as to make sense of the sharp cutoffs.

To sum up: epistemicism requires commitments to cutoffs phrased in terms of the language of fundamental physics. The sort of metasemantics that the epistemicists would require to underwrite sharp cutoffs is a significant modification of standard use-based metasemantics they rely on. This is a cost of their view.

References

Bird, Alexander, and Emma Tobin. 2024. "Natural Kinds." In *The Stanford Encyclopedia of Philosophy,* Spring 2024, edited by Edward N. Zalta and Uri Nodelman. Metaphysics Research Lab, Stanford University.

Graff, Delia. 2002. "An Anti-Epistemicist Consequence of Margin for Error Semantics for Knowledge." *Philosophy and Phenomenological Research* 64 (1): 127–142. ISSN: 00318205, accessed January 5, 2026. http://www.jstor.org/stable/3071023.

Keefe, Rosanna. 1995. *The Philosophical Quarterly (1950-)* 45 (180): 392–394. ISSN: 00318094, 14679213, accessed January 5, 2026. http://www.jstor.org/stable/2219664.

Magidor, Ofra. 2019. "Epistemicism and the Sorites Paradox." In *The Sorites Paradox,* edited by Sergi Oms and Elia Zardini. Cambridge University Press.

Sainsbury, R. M. 1996. "Concepts Without Boundaries." In *Vagueness: A Reader,* edited by Rosanna Keefe and Peter Smith, 186–205. MIT Press.

Schiffer, Stephen. 1999. "The Epistemic Theory of Vagueness." *Philosophical Perspectives* 13:481–503. https://doi.org/10.1111/0029-4624.33.s13.21.

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

Sorensen, Roy. 2001. *Vagueness and Contradiction.* Oxford University Press.

Williamson, Timothy. 1994. *Vagueness.* Routledge.

———. 2004. "Philosphical 'Intuitions' and Scepticism About Judgement." *Dialectica* 58 (1): 109–153. https://doi.org/10.1111/j.1746-8361.2004.tb00294.x.

Wright, Crispin. 1995. "The Epistemic Conception of Vagueness." *Southern Journal of Philosophy* 33 (S1): 133–160. https://doi.org/10.1111/j.2041-6962.1995.tb00767.x.