How to Resist Epistemicism

John MacFarlane

March 13, 2019

Abstract

According to epistemicists, there is a precise height which separates people who are tall from those who are not tall, though we can never know what it is. This view has struck many as preposterous, but it is harder to resist than one might think. For what seems most hard to accept about it—that vague words like 'tall' impose unknowable semantic boundaries—is also a commitment of alternative, nonclassical semantic theories. To resist epistemicism, we need two things: an argument that vague terms cannot impose unknowable semantic boundaries, and a sketch of a viable alternative—a theory of meaning that does without unknowable boundaries. I attempt to provide both.

1 INTRODUCTION

In discussions of vagueness, the term 'epistemicism' is usually used for the combination of two conceptually independent theses. The first is

> **Bivalence** If S can be used to make a literal assertion at a context c, then either S is true at c or S is false at c (that is, the negation of S is true at c).[1]

Bivalence implies that, if one were to line up all humans by order of height, the shortest human who could truly be called 'tall' in that context would be standing right next to the tallest human who could truly be called 'not tall'. Clearly, we have no way of knowing where this boundary lies. So, those who accept Bivalence are committed to a second view, which I'll call

> **Hidden Boundaries** Vague terms and concepts impose unknown and unknowable semantic boundaries.

Defending Hidden Boundaries requires giving an explanation of this epistemic limitation: if there is a fact about where exactly the boundary between the tall and the not tall lies, why are we unable to know it? Here the epistemicist typically appeals to an

> **Epistemic Explanation** We are ignorant of the locations of the hidden semantic boundaries because knowledge requires safety and hence a "margin for error." Because meaning depends on use in a complex and unsurveyable way, any belief about the exact location of the threshold could easily have been false.

The idea that our patterns of use of vague terms establishes semantic boundaries that are beyond our ken has struck many as hard to accept. But we cannot avoid the commitment to Hidden Boundaries just by dropping Bivalence and adopting a nonclassical semantics. For suppose there are truth-value gaps: sentences that are neither true nor false at a context. There will still be an unknown and unknowable semantic boundary between the shortest human who can truly be called 'tall' and the tallest human who cannot truly be called 'tall'. Or suppose truth comes in degrees between 1 (complete truth) and 0 (complete falsity). There will still be an unknown and unknowable boundary between the shortest human to whom 'tall' applies to degree 1 and the tallest human to whom it applies to a degree less than 1. These nonclassical semantics are just as committed to Hidden Boundaries as the classical, bivalent semantics.

It is this fact that makes epistemicism so hard to resist. What seems objectionable about epistemicism is not Bivalence itself, but the commitment to Hidden Boundaries. So, anti-epistemicist views that reject Bivalence but accept Hidden Boundaries lack clear motivation. Whether or not we accept Bivalence, we will need an explanation of why we cannot know where these boundaries lie. And once we have such an explanation, we might as well accept Bivalence.[2]

To reject epistemicism, then, we need two things: a compelling argument against Hidden Boundaries, and a genuine alternative—an account of vague language that doesn't posit unknowable semantic boundaries. I will attempt to provide both.

2 AN ARGUMENT AGAINST HIDDEN BOUNDARIES

The argument against Hidden Boundaries depends on two very general premises—one about speaker meaning, one about vague language—and does not rely on any specific semantic analysis. For simplicity, I'll focus on gradable adjectives like 'tall' or 'large'.

2.1 TRANSPARENCY OF MEANING INTENTIONS

Let's start with an important insight of Grice. To make a meaningful utterance, it is not enough to utter something with the intention of producing an effect in one's audience. You must intend to produce this effect through the audience's recognition of your intention:

> for x to have meant anything, not merely must it have been "uttered" with the intention of inducing a certain belief but also the utterer must have intended an "audience" to recognize the intention behind the utterance (Grice 1957, 382).

That is, you must have an intention that can only be satisfied if your hearers come to know that you have it.[3]

If you gasp at hearing some news, with the aim of getting your boss to think you didn't know the news already, you haven't *meant* by your gasp that you didn't know it already, because you didn't intend your boss to recognize your intention. But a very similar gasp, in another context, might count as a meaningful utterance—perhaps as an ironic comment to the effect that the news is hardly surprising—if it is made with an intention that requires its own recognition.

We can put the insight as follows: meaning intentions are *essentially transparent*. They aim at their own recognition. This does not mean that they are always recognized. Far from it: very often we are misunderstood, or only incompletely understood. Claire says, "Mila saw one too." One what? A praying mantis, or a centipede? We'd been talking about both. I don't know what she intended. But she did intend something, and she intended me to recognize what she intended. To ask what she meant is to ask what communicative intention she intended me to recognize.[4]

I take this Gricean insight to be fundamentally sound. That's not to say that there isn't room to worry about the precise analysis of speaker meaning that Grice gives. Many have raised doubts, in particular, about Grice's idea that a meaning intention is *inter alia* an intention to get one's interlocutor to believe something. So I'll focus on a more contemporary realization of the Gricean insight, due to Stalnaker (Stalnaker 1978, 2002, 2014). Stalnaker thinks of conversation as a rational cooperative activity governed by an evolving set of commonly accepted assumptions—the *common ground*. Assertions can be viewed as proposals to update the common ground: when an assertion is made and accepted, its content is added to the common ground. In this context, the Gricean insight can be spelled out as follows:

Transparency When one makes an assertion or other speech act, one intends it to be

shared knowledge among the parties to a conversation what update to the common ground is being proposed.[5]

Shared knowledge of the proposed update is a condition for the existence of a common ground. For, imagine what would happen if there were doubt or disagreement about how the common ground was to be updated when an assertion was accepted. Everyone would update in a different way, and there would be no common ground. (Indeed, these considerations point toward a requirement of common knowledge, not just shared knowledge, but the weaker condition is enough for our purposes here.)

In Stalnaker's framework, Transparency manifests itself as a fundamental pragmatic principle, sometimes called

Uniformity The same proposition is expressed relative to each possible world in the context set [the set of worlds that are not ruled out by the common ground]. (Stalnaker 1999, 88)

In other words, it must be common ground which proposition has been asserted.

Stalnaker argues that our interpretation of what a speaker means to assert is often guided by the requirement of Uniformity. Suppose we hear a woman talking in the next room, and I say

(1) That is either Zsa Zsa Gabor or Elizabeth Anscombe.

The common ground leaves it open that the woman is Zsa Zsa Gabor, and also that she is Elizabeth Anscombe. If she is Zsa Zsa Gabor, then the referent of 'that' in context is Gabor; if she is Elizabeth Anscombe, then the referent of 'that' in context is Anscombe. So, if (1) were being used to assert a singular proposition—if the demonstrative 'that' in (1) rigidly denoted a particular woman—then we would violate Uniformity. It wouldn't be common ground which proposition was being asserted. In order to respect Uniformity, then, we interpret the speaker as asserting a different proposition: the proposition that the woman who is talking in the next room (whoever she might be) is either Zsa Zsa Gabor or Elizabeth Anscombe.

I said I'd be mounting an argument against epistemicism. Transparency will be one premise of the argument, but by itself, Transparency is perfectly compatible with epistemicism. Suppose a speaker asserts

(2) George is tall.

According to epistemicism, there is a sharp threshold for counting as tall. One might think that our ignorance of this threshold poses an obstacle to our mutually knowing what update the speaker intends. And it would, if we took the speaker to be asserting the proposition that George's height surpasses x, for some specific height x. But there is no difficulty if—applying the lesson from Stalnaker's Anscombe/Gabor case—we take the speaker to be asserting the proposition that George's height surpasses the threshold for tallness, whatever it might be.

Indeed, our ignorance of the threshold poses no more problem for (2) than my ignorance of the exact height of the Empire State Building does for

(3) The Willis Tower is taller than the Empire State Building.

It's true that I don't know whether updating with (2) requires me to rule out worlds where George is six feet tall. But by the same token, I don't know whether updating with (3) requires me to rule out worlds where the Willis Tower is 1300 feet tall. As long as we share some general mode of presentation of the height—as the height of the Empire State Building—we can grasp the proposed update in (3). Similarly, as long as we share some general mode of presentation of the threshold for 'tall'—as the threshold for 'tall', for example!—then we can grasp the proposed update in (2).

Thus, Transparency itself poses no problem for epistemicism. Even if there is a sharp height threshold for counting as tall, knowing the meaning of 'tall' does not require knowing it.[6]

2.2 THE CONTEXTUAL FLEXIBILITY OF VAGUE LANGUAGE

You may object: "But there isn't a single threshold for being tall! The word 'tall' is contexually sensitive, and what it takes to be tall varies with context." That is certainly true. The question is whether acknowledging the contextual sensitivity of gradable adjectives poses a problem for epistemicism.

Gradable adjectives exhibit two distinct kinds of contextual sensitivity. The first and most familiar kind is relativity to a comparison class. A tree can be tall for an apple tree without being tall for a tree. A BMW may be expensive without being expensive for a BMW (Kennedy 2007, sec. 2.2). The comparison class can be made explicit, as it is in the phrase 'expensive for a BMW', but often it left implicit. When I say that Claire is tall, I mean that she is tall for a second-grader.

Does acknowledging this kind of contextual sensitivity pose any new problem for epistemicism? It's not clear why it would. We now need the threshold for tallness to vary

with a contextually determined comparison class. But that is not a problem, as long as we can find some way of picking out the threshold that will be available to all the participants in a conversation. Where before we identified the threshold as *the minimum degree of height needed to count as tall*, now we can identify it as, say, *the minimum degree of height needed to count as tall for an* F. Provided that 'tall for an F' is not *itself* contextually sensitive, Transparency can be satisfied.

One might worry that it can be hard to recognize which comparison class the speaker intends. Does the person who utters (2) mean to assert that George is tall for an American male, or that he is tall for a Californian adult male, or that he is tall for an American university professor? That may not always be clear. But that's okay: we know that communication sometimes misfires. As long as the speaker has a particular comparison class in mind and intends that the hearer recognize it, Transparency is satisfied. It would be worrisome if speakers never had a specific comparison class in mind, or if hearers could never recognize the comparison classes intended by speakers, but this does not seem to be the case.

However, in specifying a comparison class for 'tall', we have not removed all of its contextual sensitivity. In other words, 'tall for an F' is itself contextually sensitive: what it takes to be tall for an F may vary from one conversational setting to another (Fara 2000; Richard 2004; Kennedy 2007; DeRose 2008).[7]

To see this, imagine that we have a fixed reference class. We're talking about a certain specific group of apples, as depicted in Figure 1.

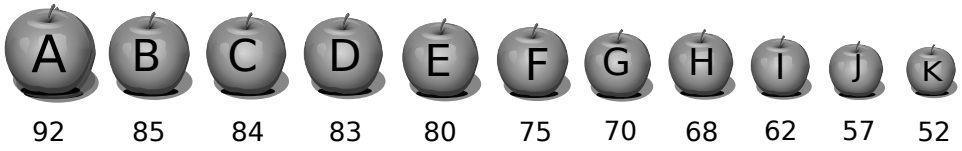

Figure 1: The apples on the table, with their diameters in millimeters (MacFarlane 2016).

It seems to me that a speaker has considerable freedom to decide what counts as large (for one of these apples). I might say, for example,

(4) Give me the large apple,

presupposing that only Apple A is large. Or I might say,

(5) Give me a few of the large apples,

presupposing that several apples are large. Which of these I do may depend on my purposes. Of course, my hearers need not go along. They can resist, saying

(6) There's more than one large apple!

in response to (4) or

(7) There's only one large apple!

in response to (5). But most of the time, my hearers will be happy to go along with the way I've chosen to talk. They will assume that I have reasons for talking the way I do, even if these reasons are not immediately apparent to them.

I suppose someone might insist that at most one of these ways of talking is correct: that it's just false, for example, to say that Apple D is large (in relation to this group), no matter what the context. I'll concede that false sentences can be used to convey useful information, and that we sometimes accommodate false claims, accepting them for the purpose of conversation. But this doesn't seem to what is happening in this case. Nothing about the conventional meaning of 'large' seems to privilege one of these ways of using 'large' over the others; rather, the meaning of 'large' is flexible enough that it can be used in any of them, drawing the distinctions among the apples that suit our purposes on different occasions.

2.3 THE ARGUMENT AGAINST HIDDEN BOUNDARIES

If there is a threshold for tallness, then, it is not determined by the comparison class alone; it can vary from one specific context to another, depending on the speaker's interests and intentions. Unlike sensitivity to a comparison class, this kind of contextual sensitivity does pose a challenge for Transparency. For although in uttering (2) I may reasonably intend that my hearers recognize that I mean to assert that George is tall for an American male, it's not plausible that I intend them to recognize that on this occasion I mean to draw the line for counting as 'tall for an American male' at, say, 192 cm.

Here the epistemicist might say: my hearers can grasp the intended threshold under the description *the threshold intended by the speaker*. Of course, they don't know whether a person 192 cm in height is over or under that threshold. But we have already granted that a threshold can be grasped under a mode of presentation that leaves its location uncertain: why not, then, as *the threshold intended by the speaker*?

The problem is that, in all but a few very unusual cases, the speaker will not have in mind any particular threshold. Even if she has a definite intention to count people

over 195 cm as tall for an American male, and people under 185 cm as not tall for an American male, she may remain undecided about people 192 cm in height. Similarly, a speaker may presuppose that there are at least a few large apples in the group, without having settled where to draw the line between the large ones and the medium ones. And if the speaker doesn't have a particular threshold in mind, it is certainly hopeless to think that the parties to the conversation can coordinate on an update specified in terms of the speaker's intended threshold.[8]

In response, the epistemicist might retreat to a more abstract specification of the threshold, like *the threshold governing this context*, or *the threshold in play in this context*. But this mode of presentation of a threshold is too thin to satisfy Transparency. Suppose I say 'That's a nice one', giving you no clues whatsoever to indicate which object I mean to be talking about. Intuitively, communication has failed; you haven't grasped my meaning intention, because you don't know which object I've demonstrated. Nobody would be tempted to object that you *do* know which object I've demonstrated, because you know it is *the object I've demonstrated in this context*, or *the referent of my use of* that *in this context*. If this were enough for you to grasp my meaning, then we would be unable to say why communication fails here.

Indeed, I think that any talk of a 'threshold of the context' is just playing with words. In what sense is there a threshold that "governs" this context, or is "in play" in it? The speaker can proceed in several different ways, none of which seem to sin against the conventional meaning of 'large'. If there were a threshold governing the context, then at most one of those continuations would be correct. But what difference would it make which it was, if we could never know? It is hard to see how a theory that purports to be an account of our linguistic knowledge could give any role to a threshold that is determined by context in ways that are forever beyond our ken.

When one is doing formal semantics, it is easy to appeal to "context" to supply values without which we won't have truth conditions. Formally, it's just a matter of writing a small '*c*'. But if the formal semantics is to help explain how we use language to speaker-mean things, then we must understand how the linguistic knowledge it codifies can help us make manifest to others what we want to assert, or more generally how we propose to update the common ground. Aspects of the formalism that don't contribute to that project should be rejected as, at best, idle wheels. Since we almost never have shared knowledge of contextually determined thresholds, and such thresholds can be relevant to meaning only to the extent that they are known, we should not posit such things in our theories of meaning.[9]

Note that this argument is directed against the *hiddenness* of the epistemicist's pu-

tative thresholds. We have said nothing about Bivalence, and so the whole argument carries over to any view that accepts Hidden Boundaries. The problem, recall, was that speakers and hearers need to coordinate on the proposed update to the common ground. This coordination would be impossible if it required coordination on a sharp threshold for 'large'. But moving to a fuzzy threshold—an assignment of real-numbered degrees of truth to various possible threshold values—would not help at all with this problem. Indeed, it would make things more difficult: since there are many more possible fuzzy thresholds than sharp thresholds, coordinating on one becomes even more difficult. If the problem is coordination, bringing in an even more fine-grained formal object does not help. Nor does it help to move to a gappy threshold—a lower and upper bound—as long as the lower and upper bounds remain hidden. If coordinating on one boundary is problematic, coordinating on two boundaries is doubly so. The problem with epistemicism, then, is its commitment to Hidden Boundaries, not its commitment to Bivalence.

3 TOWARDS AN ALTERNATIVE

As I mentioned at the outset, one of the strongest arguments for epistemicism is a *tu quoque* objection. We don't avoid unknowable semantic boundaries by adopting a non-classical semantics. If we are going to have to accept hidden boundaries anyway, we might as well stick with the simpler classical semantics. If we are going to reject hidden boundaries, then, we need a real alternative. We need a semantics for vague gradable adjectives that posits no semantic boundaries except those that, in normal communicative exchanges, are mutually known to speaker and hearer.

Here I'm not going to give a compositional semantics for gradable adjectives, but something like a "prolegomenon to any future semantics." I'll explain how we should understand the contents asserted using gradable adjectives, and how we should think of the common ground. Once this is settled, we can focus on the technical problem of giving a compositional account that generates the sorts of contents vague assertions have. I do not think there are any technical obstacles here; indeed, with a bit of reconceptualization, existing semantics for gradeable adjectives (like that of Kennedy 2007) can be ported wholesale to the new pragmatic framework. But the details will have to wait for a later occasion.

3.1 FROM THRESHOLDS TO CONSTRAINTS

Let's focus, then, on what we *do* coordinate on. Though we rarely coordinate on a specific threshold value, we routinely coordinate on *constraints*. It may be common knowl-

edge in a conversation, for example, that Apple B is the smallest apple that has been called 'large', and that Apple E is the largest apple that has been called 'not large'. This common knowledge constrains the threshold for 'large' between the size of Apple E and the size of Apple B.

As a conversation continues, the participants may narrow down these constraints. For example, if someone calls Apple C large and this is accepted, it now becomes common ground that apples at least as big as Apple C count as large, too. Indeed, sometimes assertions aim *only* to modify the constraints on the threshold, without communicating any factual information at all. If it is already common ground that Apple C is 84 mm diameter, then the point of asserting that Apple C is large may be just to establish a new upper bound for the threshold.[10] But this narrowing of the constraints should not be thought of as the reduction of uncertainty. It's not as if we're ignorant of where the threshold lies, and we're gradually narrowing the range of epistemic possibilities, reducing our uncertainty about which threshold governs our context. There is no "threshold of the context"—only constraints. So there is no fact about the actual position of the threshold to be uncertain about.

What this means is that we can no longer think of the update proposed by an assertion simply as the addition of a factual proposition to the common ground. The update can also include changes to the constraints on thresholds, where this is not to be understood as reducing uncertainty about a matter of fact. How then, should we model the common ground and the update?

3.2 THE SCOREBOARD MODEL

A natural approach would be to think of the common ground as having two components: a factual common ground—what is commonly accepted about the world—and a separate "scoreboard" (to use the term coined by Lewis 1979). The scoreboard includes settings for various nonfactual parameters: for example, the point of view that determines whether we say something is 'coming' or 'going', the standard of precision that determines whether we can say that France is hexagonal, the set of possibilities that we need not rule out in order to count as knowing something, and the constraints on thresholds for gradable adjectives.

The nonfactual scoreboard affects how we update the factual part of the common ground. So, if someone says

(8) George is coming,

we update the factual common ground by adding the proposition that George is approaching the location the scoreboard identifies as the "point of view." In general, we do this by updating our view of the facts in whatever way is required to make the assertion true, as interpreted in light of the nonfactual scoreboard.

If the nonfactual scoreboard provides only a constraint—as with thresholds—then the factual common ground will be updated by adding whatever is required to make the assertion true on every way of satisfying the constraints.[11] For example, if the scoreboard currently limits the threshold to between the size of Apple B and Apple E, then an assertion of

(9) George saw a large apple

will add to the factual common ground that George saw an apple at least as large as Apple B.

The main mechanism for making changes to the nonfactual scoreboard is what Lewis calls *accommodation*:

> If at time t something is said that requires component s_n of conversational score to have a value in the range r if what is said is to be true, or otherwise acceptable; and if s_n does not have a value in the range r just before t; and if such-and-such further conditions hold; then at t the score-component s_n takes some value in the range r. (Lewis 1979, 347).

Thus, for example, if the scoreboard currently calls for high standards of precision, and someone who is not ignorant of geography asserts

(10) France is hexagonal,

we can "accommodate" this assertion, allowing it to be true, by relaxing the scoreboard's setting for standards of precision. Similarly, if we know that Apple D is 83 mm diameter, and the scoreboard currently constrains the threshold for 'large (for an apple)' to between 92 and 75 mm, we can accommodate an assertion of

(11) Apple D is large

by revising the upper bound for the threshold to 83 mm. Notice, though, that accommodation is only called for when there is no way to interpret the assertion as true without changing the nonfactual scoreboard. So, we have a

Recipe for updates

a) Figure out what the world must be like if the sentence is to be true on all thresholds compatible with the constraints.

b) Ask: Is this proposition compatible with the factual common ground?

   - If yes: update the factual common ground.
   - If no: update the nonfactual constraints (and perhaps the factual common ground as well).

Though this is a natural and familiar picture, I believe it is fatally flawed. The factual and nonfactual components of the common ground cannot be separated in this way.

To see why, suppose that Apple C is the smallest apple that has been called 'large' in this context, and Apple F is the largest apple that has been called 'not large', so that the nonfactual scoreboard constrains the threshold for 'large' to between the size of Apple C and the size of Apple F. Now, suppose we introduce a new Apple, X. It appears to be about the same size as Apple C—maybe bigger, maybe smaller—one can't tell by looking. Someone asserts:

(12) Apple X is large, too.

What is the proposed update to the common ground?

The view we have just sketched makes a prediction, but I think the prediction is wrong. Remember, accommodation is only called for when there is no way to interpret the assertion as true without adjustments to the nonfactual scoreboard. Here that condition is not met. For (12) can be true on all admissible settings for the threshold, provided Apple X is at least as large as Apple C. And it is compatible with the factual common ground that X is at least as large as C. So, that is the recommended update. We add to our factual common ground the proposition that

(13) Apple X is at least as large as Apple C

and we do not adjust the constraints on thresholds.

But this, I submit, is the wrong update. One can assert (12) in this context without ruling out the possibility that Apple X is a bit smaller than Apple C. All one is committing oneself to is that the size of Apple X exceeds the threshold for largeness. If one later discovers that Apple X is a bit smaller than Apple C, then one need not retract (12).

The problem cannot be fixed by relaxing the recipe and allowing changes to the nonfactual scoreboard even in cases where they are not needed for a consistent update to the factual common ground. For, what change to the constraints on thresholds for 'large'

would be called for in response to (12)? One wants to say: it depends on the actual size of Apple X. If Apple X is actually smaller than Apple C, then a lower upper bound on the threshold is called for; if not, not. But the update to the common ground cannot be conditional on something that isn't common ground, like the size of Apple C, or Transparency will be violated.

The point can be made even more simply by considering a disjunctive assertion:

(14) Either every apple over 82 mm diameter is large, or Apple X is not over 82 mm diameter.

What update to the common ground is required if (14) is accepted? Not an update to the nonfactual scoreboard, because (14) is compatible with the current settings of the scoreboard. And not an update to the factual common ground, because (14) is compatible with the current factual common ground. So it seems (14) does not correspond to any update to a bifurcated common ground. Yet (14) is not a trivial assertion. One might have grounds for rejecting it, and accepting it constrains what one says in the future.

3.3 THE INTEGRATED MODEL

What is the right update in response to an assertion of (12)? It is neither a factual update (13), nor an update to the constraints on thresholds, nor a combination of these. Asserting (12) requires no change at all either to our factual information or to the range of thresholds, considered separately. What it rules out are certain *combinations* of thresholds and sizes for Apple X: combinations that assign Apple X a size that is lower than the threshold. But this kind of update cannot be represented in our bifurcated common ground.

We can represent it if, instead of thinking of the common ground as a pair consisting of a set of worlds and a set of thresholds, we think of it as a set of pairs of a world and a threshold.[12] Asserting (12) has the effect of ruling out every pair of a world w and threshold t such that the size of Apple X in w is less than t.

This formalization also solves another problem with the bifurcated representation. In general we won't be able to say what the constraints on thresholds are in, say, millimeters. We know that the threshold is constrained between the sizes of Apple C and and Apple F, but we don't know the exact sizes of these apples. So it was never really the case that the constraints on thresholds could be specified independently of the constraints on worlds. Representing the common ground as a set of world/threshold pairs allows us to make that dependence explicit.

This formal model of the dynamics of vague language can be found in Chris Barker's pioneering work (Barker 2002, cf. 2009, 2013). However, Barker spoils the insight by thinking of the thresholds not as a genuinely non-factual component of the common ground, but as a special kind of fact—a fact about the "discourse" (Barker 2002, 5):

> Just as we can peer into a world w and determine whether it is raining or whether Abby is a doctor in that world, we will peer into a discourse situation d in order to see whether a certain index refers to a man or a woman, or what the height cutoff for counting as tall happens to be in d (Barker 2009, 257).[13]

This leads him to conceive of assertions as reducing our "ignorance" or "uncertainty" (Barker 2002, 3–4, 9) about "the" threshold governing the discourse. The ideology of Hidden Boundaries has returned. But *just* the ideology. The hidden thresholds play no role in Barker's account of how vague assertions update the common ground. Grasping the update requires only knowing the constraints on thresholds, not the hidden "actual" threshold.[14]

Still, I must protest that the ideology is problematic. There *is* no feature of the discourse situation that could be picked out as "the height cutoff for counting as tall" in this situation. The participants are leaving themselves the flexibility to narrow down the range of cutoff points as the conversation progresses. It is possible that the conversation may eventually require fixing on a determinate, known cutoff point (for example, if its purposes require classifying everyone either as tall or as not tall). But even in this unusual case, the participants cannot be described as *discovering* a threshold that governed their discourse situation all along. Rather, they are *deciding* on a threshold.

Consider the apples again. Suppose, as before, that we all know that Apple C is the smallest apple that has been called 'large', and that Apple F is the largest that has been called 'not large'. And we know the diameter of Apple E in millimeters. It seems to me that it would be very odd to express our ambivalence about whether Apple E is large using the language of epistemic uncertainty:

(15) Apple E might be large.
(16) I can't tell whether Apple E is large.
(17) Apple E is probably large.

We would more naturally use the language of decision:

(18) We could count Apple E as large.
(19) I can't decide whether to count Apple E as large.

(20) We should probably count Apple E as large.

By contrast, when we're drawing a comparison with an unknown size, it seems perfectly reasonable to use the epistemic language, and incoherent to use the language of decision. Suppose we can see a baseball, but we only have inexact perceptual knowledge of its size. Then it would be fine to say

(21) Apple E might be as large as that baseball.
(22) I can't tell whether Apple E is as large as that baseball.
(23) Apple E is probably as large as that baseball.

but quite odd to say

(24) We could count Apple E as as large as that baseball.
(25) I can't decide whether to count Apple E as as large as that baseball.
(26) We should probably count Apple E as as large as that baseball.[15]

Epistemicists have focused on explaining why we can't have *knowledge* of the exact location of a threshold. If knowledge is a norm for belief, this might also explain why we don't have beliefs about the exact location of a threshold. But the examples above show that more is needed. Although we can't *know* whether Apple E is bigger than the baseball, it may be perfectly appropriate to think that it is probably bigger, or no more likely than not to be bigger. If there is a fact of the matter about where the threshold for 'large' lies, why can't we have intermediate credences about whether Apple E's size exceeds it? And why are we willing to use the language of decision in one case but not the other? Some explanation is needed, and I don't see much promise of an epistemic explanation here.

What we need, rather, is a way of understanding the threshold component of the pairs in our integrated common ground as relating to *decision* in the way that the world component relates to belief.

3.4 EXPRESSIVISM

This way of formulating the problem points towards a solution. Instead of understanding the threshold component of our pairs as representing a kind of fact—a fact about the "discourse situation"—we can understand it as representing a kind of plan or intention: the plan to use such-and-such a threshold in interpreting 'large' in this conversational context.

Of course, we rarely decide to use a precise threshold. We *could* do so—I think nothing about the conventional meaning of gradable adjectives prevents us from doing so—but our purposes almost never call for it.[16] Usually our plans are partial: for example, we decide that anything over a certain size will count as 'large'. But, as Allan Gibbard has suggested (Gibbard 2003), we can represent a partially undecided plan as a set of fully determinate plans, or *hyperplans*, just as we can represent the content of a belief as a set of fully determinate belief contents, or possible worlds. If I have decided to count every apple over 85 mm diameter as 'large', then I have ruled out hyperplans that draw the line at 86 mm, 87 mm, and so on. The content of my plan can be represented as the set of all hyperplans compatible with it (including ones that draw the line at 85 mm, at 84 mm, and so on).

However, many of our mental states are neither pure doxastic states nor pure planning states. If P is the content of a pure doxastic state and Q is the content of a pure planning state, then accepting their disjunction is presumably neither being in a pure doxastic state nor being in a pure planning state. To handle such cases, Gibbard argues, we need to represent the contents of assertions and attitudes as sets of (world, hyperplan) *pairs*. And that is how I propose to represent the contents of assertions and the common ground. A Gibbardian hyperplan, of course, will include much more than instructions for how to draw a threshold for 'large'. It will include contingency plans for all kinds of things, like what you should say if you receive a call offering to appoint you Ambassador to France. But for our purposes, we can ignore everything that doesn't matter for our purposes, and represent the hyperplan as a threshold setting.

Let's see how this helps with the case that proved difficult for the bifurcated common ground. To review: It's common ground that Apple C is the smallest apple that has been called 'large' in this conversation, and we know it's 84 mm diameter. Apple F is the largest apple that has been called 'not large', and we know it's 75 mm diameter. Apple X is about the same size as Apple C, but may be slightly bigger or smaller. Simplifying a bit, we can say that the only *factual* uncertainty in the common ground concerns Apple X's size, which might be anywhere from 83–85 mm. Also for simplicity, we'll ignore the lower bound of the threshold for 'large', and consider only the upper bound. Then the common ground can be represented as a set of pairs consisting of a size s for Apple X and a threshold upper bound u. Initially, these are independent, so the common ground contains every combination:

(27) $\{\langle s, u\rangle \mid 83 \leq s \leq 85, 75 < u \leq 84\}$.

The content of the assertion that Apple X is large (12) is the set

(28) $\{\langle s, u\rangle \mid s \geq u\}$.

So the update consists in removing all pairs from the common ground in which Apple X's size is not greater than or equal to the upper bound of the threshold for 'large'. We get this update by taking the intersection of (27) and (28):

(29) $\{\langle s, u\rangle \mid s \geq u, 83 \leq s \leq 85, 75 < u \leq 84\}$.

This update does not rule out any possibilities for the size of Apple X, or any threshold upper bounds: in particular, (29) does not accept that Apple X is larger than Apple C, so we avoid the problematic prediction of the bifurcated view. But the update does rule out certain *combinations* of sizes and threshold upper bounds. For example, the old common ground contained the pair $\langle 83, 84\rangle$, but the new one does not. So the update is not trivial.

On this picture, the common ground is not, as in Stalnaker, a state of common belief, but rather a common doxastic-planning state, defined in terms of joint planning and common belief:

- For it to be common ground that any apple more than 84 mm in diameter is large is for us to have a joint plan to use 'large' (for an apple) in a certain way.
- For it to be common ground that Apple X is 84 mm in diameter is for us to be in a joint doxastic state of common belief about what is accepted.
- For it to be common ground that Apple X is large is for us to be in a hybrid state which is incompatible with, e.g., commonly believing that Apple X is accepted to be 83 mm diameter and jointly planning to count only apples more than 84 mm diameter as large.

When we make assertions using vague words like 'large', we are typically constraining both our view of the world (the doxastic possibilities) and our plans for using the word (the practical possibilities), with the aim of getting others to coordinate on both a common picture of the world and a common plan for using words. But the constraints we impose are entangled, and cannot be separated into separate factual and nonfactual updates, as the bifurcated view assumes.

4 CONCLUSION

I have argued that what is objectionable about epistemicism is its commitment to Hidden Boundaries. Given the contextual flexibility of vague words, which can be used in very different ways even when a comparison class is fixed, Hidden Boundaries is incompatible with the transparency of meaning intentions. The epistemicist cannot explain how it can be mutually known what update is being proposed when someone says, for example, that a certain apple is large.

If the problem with epistemicism is Hidden Boundaries, it does not help at all to embrace alternatives that reject Bivalence while still being committed to Hidden Boundaries. To reject Hidden Boundaries, I have argued, we must reject the view that context determines a threshold (even a "fuzzy" one) for vague gradable adjectives. We must replace the standard conception of the common ground as a body of factual propositions (or a set of possible worlds) with a conception that makes room for nonfactual constraints on thresholds for vague gradable adjectives. The range of threshold values left open by these constraints does not represent uncertainty. There is simply no fact of the matter where, within these constraints, the threshold lies. That is why it is inappropriate to use the language of subjective uncertainty, and appropriate to use the language of decision, in borderline cases.

I have argued further that the nonfactual component of the common ground cannot be separated from the factual component. Instead, we need an integrated common ground, which can be thought of as a joint hybrid doxastic and planning state. The same model will work for the individual mental states whose contents are attributable using vague language. To resist epistemicism, then, one ought to embrace a form of Gibbardian plan expressivism.

That is not to say that one must be an expressivist about normative language. I think there are some good objections to normative expressivism that do not carry over to the kind of expressivist advocated here. We can thank Gibbard for the philosophical tools he has developed, while declining to use them for the purpose he intended. Note that even a realist about the *better than* relation faces the problem of how to understand the positive form 'good', and should prefer an expressivist approach to the threshold rather than an epistemicist one, for the reasons given above.

We noted at the outset that the epistemicist takes Hidden Boundaries to be a consequence of Bivalence. Does the expressivist, then, reject Bivalence?

At the outset, we stated Bivalence as the view that every sentence apt for making a literal assertion at a context c is either true or false at c. On the expressivist view, a con-

text does not determine thresholds that would settle extensions for gradable adjectives: to assign truth values to vague sentences, we need to specify not just a context but a hyperplan. Thus an expressivist must reject the idea that vague sentences have truth values relative to contexts, and cannot accept Bivalence as we initially formulated it.

But there are other formulations of bivalence that the expressivist may be able to accept. Following Williamson (1994, 188), we can define monadic truth and falsity predicates, applicable to utterances, as follows:

(30) If an utterance u says that ϕ, then u is true iff ϕ and false iff $\neg\phi$.

As Williamson acknowledges, the predicates 'true' and 'false' in (30) are vague: if Sam is a borderline case of a tall man, then an utterance of 'Sam is a tall man' will be a borderline case of a truth.[17] Using these predicates, we can formulate bivalence as the schema

(31) If u is an utterance that says that ϕ, then u is either true or false.

Given (30), (31) is equivalent to a restricted version of the principle of excluded middle:

(32) If u is an utterance that says that ϕ, then ϕ or $\neg\phi$.

Nothing about the expressivist view rules out accepting excluded middle. An expressivist, then, can accept classical logic and even bivalence, as formulated in (31), without embracing Hidden Boundaries and the epistemic view.[18,19]

NOTES

1. Williamson (1994, 187) formulates bivalence differently, as the thesis that every utterance that expresses a proposition is either true or false. Williamson's formulation says nothing about sentence/context pairs that do not correspond to utterances, but otherwise agrees with this one, on the assumption that a sentence S can be used to make a literal assertion at a context c just in case an utterance of S at c would express a proposition, and that such an utterance would be true (or false) just in case S is true (or false) at c.

2. This is the core argument of Williamson (1994), which convinced many people to take epistemicism seriously.

3. Arguably, something stronger is required: that your intention is only satisfied if it becomes *common knowledge* that you have it. (Your hearer knows that you have it, you

know that they know this, they know that you know that they know it, and so on.) For our purposes here, we won't need anything that strong (for discussion, see Schiffer 1972; Grice 1969, 156, 1989, 302–3).

4. That's what is so strange about Humpty Dumpty (in Lewis Carroll's *Through the Looking Glass*). He claims to be able to mean *that's a nice knockdown argument* by 'There's glory for you', but can he really have the the intention that we recognize that this is what he wants to get across? Normally one cannot intend to do something one thinks is impossible. But Humpty is a strange egg.

5. Stalnaker (2002, 704) explicitly connects his idea of a common ground with Grice's analysis of speaker meaning. "One thing, according to Grice, that is distinctive about speaker meaning, as contrasted with other ways of getting people to believe something, is a kind of openness or transparency of the action: when speakers mean things, they act with the expectation that their intentions to communicate are mutually recognized. This idea leads naturally to a notion of common ground...."

6. How might one give a semantics for 'tall' that posits a sharp threshold but doesn't imply that knowledge of that threshold is required for knowledge of meaning? One might simply give the meaning homophonically: $[\![\textit{tall}]\!] = \{x : x \text{ is tall}\}$ (cf. Evans and McDowell 1976, xi; Sainsbury 1996, sec. VII). Speakers can know that 'tall' has this extension without knowing how tall something has to be, in millimeters, to count as tall.

   One might object that this approach leaves us without a good account of the meaning relations between 'tall', 'taller', and 'tallest'. To handle that, one might switch to a degree semantics, on which the semantic value of 'tall' is a function from objects to abstract degrees of tallness (see e.g. Kennedy 2007). On this approach, the extension of the positive form of the adjective 'tall' would be the set of objects mapped by the degree function to a degree exceeding some threshold. But nothing in this machinery prevents us from stating the threshold homophonically, as "the threshold for tallness" or "the minimum degree of height needed to count as tall."

7. Kennedy's degree semantics makes room for this additional dimension of contextual freedom by supposing that context provides a function which, applied to any degree function restricted to a comparison class, fixes a threshold (Kennedy 2007, 17).

8. Might we say that the speaker intends a *vague* threshold, and that hearers can recognize which vague threshold is intended? One should not assume that vagueness affects only sentences, and not our beliefs, intentions, and other mental states. But this just raises

the question what is it to intend a vague threshold, or to believe that a speaker intends a particular vague threshold. The proposal at the end of this paper may be regarded as an attempt to answer this question.

9. Jeffrey King puts the point succinctly: "Not much use having a value assigned to a supplementive if it plays no role in communication!" (King 2014, 110; cf. MacFarlane 2016, 261–2). This is a special case of a general point made by James Higginbotham: "Statements of truth conditions that go beyond these bounds [what is known in common] are irrelevant to understanding, resting as it does on common knowledge, and so irrelevant to meaning as well" (Higginbotham 1991, 10).

10. Barker (2002, 2) imagines asserting 'Feynman is tall' in a context where everyone knows exactly how tall Feynman is. The point of this assertion, Barker says, "would be nothing more than to communicate something about how to use a certain word appropriately." I will say more about Barker's view in what follows.

11. Thus, the view under consideration is a version of supervaluationism, in which the set of admissible valuations can vary with context.

12. In the end, we'll need a threshold *function* that assigns a threshold to any domain-constrained degree function (Kennedy 2007), but if we're just thinking about 'large for an apple', we can use a threshold value to keep things simple.

13. The same idea is reiterated with some hesitation in Barker (2013, 245–6). Note that in Barker (2002, 6), d is not even formally independent of w; it is determined as a function of the world of evaluation.

14. One way to get to a picture like Barker's would be to start with the idea that there are determinate but unknown facts about threshold values in a context, and then note that Stalnaker's Uniformity constraint calls for diagonalization. The proposition expressed by 'George is tall' would then be something like: $(t = 200 \wedge h \geq 200) \vee (t = 201 \wedge h \geq 201) \vee \ldots \vee (t = 220 \wedge h \geq 220)$, where t is the height threshold and h is George's height. The problem, I think, is that *grasping* this diagonal proposition requires having a conception of *the threshold governing the context*. This, as I have argued, is a fiction.

    Moreover, if we think of the threshold as an aspect of the world, then our ambivalence about the classification of borderline cases should be understood as ordinary uncertainty, and this has some counterintuitive results, as first noted by Schiffer (2003, chap. 5). If we are ambivalent about whether to classify Sam as a tall man, as a bald man, as a smart man, and as a funny man, then we'll also be ambivalent about whether to classify

him as a tall, bald, smart and funny man. But in order to count as a tall, bald, smart and funny man, Sam must exceed the thresholds for all four gradeable adjectives. Assuming the positions of the thresholds are reasonably independent of each other, the probability that he exceeds all four should be *much* smaller than the probability that he exceeds any of them singly. We should be pretty confident, then, that Sam is *not* a tall, bald, smart and funny man. But in fact, Schiffer observes, we are about as ambivalent about the conjunction as we are about the conjuncts.

15. I don't want to say that (24), (25), and (26) could never be appropriate. They can be appropriate when there is more than one reasonable way to map the physical dimensions of the objects to degrees of largeness. For example, if Apple E is wider than it is tall, and it is wider than the baseball but also shorter, we might regard it as to some degree a matter of decision how we should apply 'larger'. But this is just a case where the comparative 'larger' is itself vague, so it supports the distinction I am making here. To get clear judgments in (15–26), assume that both Apple E and the baseball are perfectly spherical.

16. For discussion, see MacFarlane (2016, sec. 7).

17. "We must accept that our attributions of truth and falsity, like just about all all other utterances, have some element of vagueness" (Williamson 1994, 192).

18. If you are puzzled about how one could accept (31) without accepting Hidden Boundaries, note that 'true' and 'false' as they are used in this principle are plan-dependent (unlike the semanticist's 'true at c'). If one plans to use 'large' a certain way—say, to apply only to apples over 83 mm—this plan will also affect one's use of the monadic predicates 'true' and 'false'. Thus, a common ground that accepts that an utterance u is either true or false may fail to accept either that u is true or that it is false: u may count as 'true' on some possible ways of firming up our plans, and false on others.

19. This paper derives from a talk given at UNC Chapel Hill on November 7, 2017. I am grateful for support from a fellowship at the Paris Institute for Advanced Studies, with the financial support of the French State managed by the Agence Nationale de la Recherche, programme "Investissements d'avenir", (ANR-11-LABX-0027-01 Labex RFIEA+), and the Fondation Maison des Sciences de l'Homme. Thanks also to Francois Recanati, Paul Egré, Chris Barker, Michael Rieppel, my Spring 2018 graduate seminar at Berkeley, and colloquium audiences in Chicago, Konstanz, Paris, Berkeley, Stanford, and Syracuse.

REFERENCES

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

———. 2009. "Clarity and the Grammar of Skepticism." *Mind and Language* 24: 253–73.

———. 2013. "Negotiating Taste." *Inquiry* 56: 240–57.

DeRose, Keith. 2008. "Gradable Adjectives: A Defence of Pluralism." *Australasian Journal of Philosophy* 86: 141–60.

Evans, Gareth, and John McDowell, eds. 1976. *Truth and Meaning*. Oxford: Oxford University Press.

Fara, Delia Graff. 2000. "Shifting Sands: An Interest-Relative Theory of Vagueness." *Philosophical Topics* 28: 45–81.

Gibbard, Allan. 2003. *Thinking How to Live*. Cambridge, MA: Harvard University Press.

Grice, H. P. 1957. "Meaning." *The Philosophical Review* 66: 377–88.

———. 1969. "Utterer's Meaning and Intention." *The Philosophical Review* 78: 147–77.

———. 1989. "Meaning Revisited." In *Studies in the Way of Words*, 283–303. Cambridge: Harvard University Press.

Higginbotham, James. 1991. "Truth and Understanding." *Philosophical Studies* 65: 3–16.

Kennedy, Christopher. 2007. "Vagueness and Grammar: The Semantics of Relative and Absolute Gradable Adjectives." *Linguistics and Philosophy* 30: 1–45.

King, Jeffrey C. 2014. "The Metasemantics of Contextual Sensitivity." In *Metasemantics*, edited by Alexei Burgess and Brett Sherman, 97–118. Oxford: Oxford University Press.

Lewis, David. 1979. "Scorekeeping in a Language Game." *Journal of Philosophical Logic* 8: 339–59.

MacFarlane, John. 2016. "Vagueness as Indecision." *Proceedings of the Aristotelin Society* s.v. 90: 255–83.

Richard, Mark. 2004. "Contextualism and Relativism." *Philosophical Studies* 119: 215–42.

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

Schiffer, Stephen. 1972. *Meaning*. Oxford: Oxford University Press.

———. 2003. *The Things We Mean*. Oxford: Oxford University Press.
Stalnaker, Robert. 1978. "Assertion." In *Syntax and Semantics, Vol. 9: Pragmatics*, edited by P. Cole. New York: Academic Press.
———. 1999. *Context and Content: Essays on Intentionality in Speech and Thought*. Oxford: Oxford University Press.
———. 2002. "Common Ground." *Linguistics and Philosophy* 25: 701–21.
———. 2014. *Context*. Oxford: Oxford University Press.
Williamson, Timothy. 1994. *Vagueness*. London: Routledge.