VAGUENESS AS INDECISION

JOHN MACFARLANE

This paper motivates and explores an expressivist theory of vagueness, modeled on Allan Gibbard's (2003) normative expressivism. It shows how Chris Kennedy's (2007) semantics for gradable adjectives can be adjusted to fit into a theory on Gibbardian lines, where assertions constrain not just possible worlds but plans for action. Vagueness, on this account, is *literally* indecision about where to draw lines. It is argued that the distinctive phenomena of vagueness, such as the intuition of tolerance, can be explained in terms of practical constraints on plans, and that the expressivist view captures what is right about several contending theories of vagueness.

I

Introduction. You've just started work in an apple sorting plant. Your job is to put all the large apples in one crate, where they'll go to the supermarket, and all the small and medium ones in another, where they'll be made into apple juice. As the apples go by, your supervisor says, 'That's a large one. So is that. That one's a medium one, and this is a small one.' You soon start anticipating her in your own thought, making your own classifications: *there goes another large one.*

What do you get from this training? It can't be wrong to say that you learn which apples are large, medium, and small. But it does seem wrong to say that learning this is learning something about the *sizes* of the apples. You knew the apples' sizes already: they were in plain view even before the supervisor spoke. We can even imagine that they are on a conveyor belt marked with a uniform grid of one-millimeter squares, so you can see at a glance how big each is. In learning which apples are large, medium, and small, you learn not how big they are, but how large an apple must be in order to *count as* large, medium, or small in this context.

1

In learning this, you learn something about how your supervisor, and presumably the other workers in the plant, use the words 'large', 'medium', and 'small'. But her assertions are not intended merely to carry information about her usage. Agreeing with her requires more than just recognizing that she uses 'large' in a certain way; it requires constraining your own usage to conform to the example she has set. Constraining your own usage is a matter not of ruling out factual possibilities, but of adopting a plan to use the word 'large' and the concept *large* in a certain way, applying them to any apples over a certain size. Accordingly, your supervisor's assertion can be thought of as a proposal to adopt a joint plan, and your agreement as acceptance of this plan. This, at any rate, is the thought I want to explore. If it is correct, then the proper theory of meaning for vague language is *expressivist*.

The kind of expressivism I want to use as my model is due to Gibbard (2003). Gibbard's central idea is that simple, all-things-considered deliberative ought claims, like

(1) I ought to pack now,

express *plans* rather than factual beliefs. Judging that one ought to pack now *just is* having a plan to pack. For the normative realist Gibbard opposes, judging that one ought to pack is recognizing the truth of some proposition. The question then arises why recognizing the truth of this proposition should motivate one to plan to pack. Gibbard proposes that we skip the middle step and say that (1) directly expresses a planning state of mind.

I want to explore a similar move with

(2) That is a large apple.

To judge (2), in a context where it is taken for granted that the demonstrated apple is 84mm across, is just to plan to apply the concept *large* (and the predicate 'large') to apples 84mm in diameter and larger. To assert (2) is to express this plan, and to propose it for joint adoption. Neither judging (2) nor accepting an assertion of (2) requires recognizing the truth of a proposition that it expresses.

If we accept Gibbard's view that normative statements express plans, then this view implies that (2) is covertly normative. But officially

we can remain neutral on this question. My case for an expressivist account of vague language will not assume that Gibbard's account of normative language is correct.

Although the approach I will recommend is in many respects radically different from standard approaches to vagueness, it partially vindicates some of their central claims. After setting out the expressivist view, I will argue briefly that it captures what is right in contextualism, nihilism, supervaluationism, and epistemicism, while avoiding what is implausible in them.

<div align="center">II</div>

Degree semantics for gradable adjectives. The word 'large' is a gradable adjective, so we would like a theory that can fully explain its compositional behavior. The root 'large' occurs in comparative and superlative constructions ('larger', 'largest'), as well as in a positive construction ('large'). There are some obvious entailments here: for example, 'x is large' and 'y is larger than x' entail 'y is large.' The positive form can be modified with a domain restriction, as in 'large for an apple.' Something can be large for an apple without being large. We can also say that one thing is large 'compared to' another, even if neither is large. We would like a compositional semantics for 'large' to explain all of these things.

A standard approach to this problem takes the semantic values of gradable adjectives to be functions from objects to degrees on a totally ordered scale (I will follow the version in Kennedy 2007, who cites a number of antecedents). Thus, for example, the semantic value of 'large' is a function (call it 'LargenessOf') that maps objects to degrees of largeness. Given this root meaning, it is easy to see how the comparative and superlative forms work. x is 'larger than' y just in case LargenessOf(x) > LargenessOf(y). x is 'largest' if for all y in the relevant domain, LargenessOf(x) > LargenessOf(y).

But what about the positive form 'large'? On Kennedy's view, x is in the extension of 'large' just in case LargenessOf(x) ≥ s(LargenessOf) (Kennedy 2007, p. 17). s is a function that takes the semantic value of 'large' and yields a *delineation*: a minimal degree of height that counts as large. (What fixes s? We'll come back to that.)

How can this theory account for the fact that we can call an apple

or an ant 'large'? When we call an apple 'large,' we generally mean that it is 'large for an apple,' though the qualification usually doesn't need to be made explicit. The role of 'for an apple,' on Kennedy's view, is to restrict the domain of the function expressed by 'large.' Where LargenessOf is a function that maps everything (or at least everything with a size) to a degree of largeness, we can use the notation 'LargenessOf$|_{\text{apple}}$' for the restriction of that function to the domain of apples. LargenessOf$|_{\text{apple}}$ is a partial function, and it is different from LargenessOf, and also from LargenessOf$|_{\text{Cortland-apple}}$, which has an even more restricted domain. So the s function, which determines a threshold degree given a function from objects to degrees, can map LargenessOf, LargenessOf$|_{\text{apple}}$, and LargenessOf$|_{\text{Cortland-apple}}$ to different thresholds on the same largeness scale. This explains how an object can be correctly characterized as 'large for an apple, but not large for a Cortland apple,' or 'small, but large for an ant.'

Kennedy explains the phrase 'x is large compared to y' by positing that 'compared to' introduces a special context where the domain includes just x and y. This explains why 'this ant is large compared to that one, but neither is large' is not contradictory.

It is commonly held that 'for an apple' contributes a 'comparison class' that affects the interpretation of 'large.' Kennedy can agree, if by 'comparison class' we understand 'domain restriction.' But this view is often combined with two others that, as Kennedy shows, are mistaken. First, it is often assumed that when 'large' is combined with a common noun, as in 'large apple,' the comparison class or domain restriction is given by the noun. Kennedy gives the following example to show that this is false:

(3) Kyle's car is an expensive BMW, though it's not expensive for a BMW. In fact, it's the least expensive model they make. (Kennedy 2007, p. 11; see also DeRose 2008)

Second, it is often assumed that the comparison class *completes* the meaning of the gradable adjective, so that once we know the relevant comparison class, we have what we need to determine what counts as 'large.' On Kennedy's view, the comparison class (or, as he thinks of it, domain restriction) only affects the *argument* of the s function. Since different s functions could map the same argument

to different thresholds, it is wrong to think that a comparison class by itself fixes a delineation. The same comparison class could yield different delineations in different contexts (Kyburg and Morreau 2000; DeRose 2008; Richard 2008; Shapiro 2008). In one context, an apple might need to be 87mm across to be 'large for an apple,' in another context, it might need to be 84mm across.

To sum up Kennedy's view, the extension for a particular use of the positive form of a gradable adjective depends on the three factors:

a) the semantic value of the adjective, which determines a mapping from things to degrees on a scale;

b) an explicit or implicit domain restriction;

c) a function (s) that takes a domain-restricted function from things to degrees, and yields a delineation—a 'cut-off' point on the scale for the applicability of the positive form.

There is just one question remaining. What fixes s? If our compositional semantics is going to spit out a condition for each sentence to be true at a context, then s must be fixed somehow by the context, and that is what Kennedy assumes:

> But exactly what function is s? The answer I will pursue here is that s is a context-sensitive function that chooses a standard of comparison in such a way as to ensure that the objects that the positive form is true of 'stand out' in the context of utterance, relative to the kind of measurement that the adjective encodes. (Kennedy 2007, p. 17)

He assumes further than s will determine a sharp cutoff point, dividing objects into two classes, and that its precise location will be unknown:

> Following Williamson (1992, 1994), we may assume that the reason why we cannot say exactly which degree is the one that determines whether an object stands out relative to g is because of epistemic uncertainty about the precise location of $s(g)$ in the context; this will also account for borderline cases. (Kennedy 2007, p. 19)

As I will argue in the next section, these assumptions about how the *s* function must be determined yield an untenable picture of how vague language works. The real problem is not Kennedy's embrace of a sharp cutoff point, but his more fundamental commitment to deriving conditions for a sentence to be *true at a context* from his semantic machinery. So the needed repair is, in a way, quite radical. On the other hand, it will keep in place all of the details of Kennedy's compositional semantics, excepting the story about how *s* is determined. And that is good, because Kennedy's degree semantics is the best approach we have to explaining the compositional behavior of gradable adjectives.

<p style="text-align:center">III</p>

Motivating expressivism. Imagine that we have on the table eleven apples of different diameters (Fig. 1):

| A	B	C	D	E	F	G	H	I	J	K
92	85	84	83	80	75	70	68	62	57	52

Figure 1: Eleven apples, with diameters in mm.

If I say, 'give me the large one,' you'll give me Apple A. But if I say, 'give me the four large ones,' you'll give me A, B, C, and D. If I say, 'there are no small apples here—could you give me a few of the medium ones?', you'll give me I, J, and K. But if I say, 'give me two of the small ones,' you'll give me J and K. In other respects, aside from the requests I make, the contexts may be very much the same. This illustrates something we already noted above: that specifying a domain or comparison class—here, the eleven apples A–K—is not enough to determine a particular delineation for 'large.' But it also shows something else. To a very large extent, the way the delineation is fixed is *up to the speaker*. Before I make my request, you would regard either 'give me the large one' or 'give me the four large ones' as legitimate requests, though the first presupposes that there is just one large apple in the domain, and the other that there are exactly four. Neither of these requests can be dismissed on the

grounds that it tries to put the boundary line somewhere it can't go. So it seems that speakers have considerable leeway to move the boundaries between 'small,' 'medium,' and 'large' as they find useful, and hearers are able to follow them (a point well made by Kyburg and Morreau 2000).

This kind of flexibility is not unique to gradable adjectives. We have plenty of other flexible expressions, whose extensions are to a great extent up to the speaker to determine. The most obvious examples are bare demonstratives like 'this' and 'that.' In principle, I can use 'that' to refer to any object. But with this freedom comes great responsibility. I must provide my hearers with enough cues to enable them to associate my use of 'this' with the same object I do, or communication will fail. Sometimes this doesn't require anything extra, because it is mutually known that one object is salient, so that it will be assumed to be the referent in the absence of cues to the contrary. Other times it requires pointing. And in some cases simply pointing isn't enough. But in every case, we're obliged to do whatever is required to get our hearers to associate the same object with the demonstrative that we do. If we fail to do this, it will be sheer luck if they understand us.

Of course, the hearer can always identify the object the speaker has in mind as 'the object the speaker is referring to with her use of the word "that".' But obviously, the availability of *this* description isn't enough to satisfy the requirement. For communication to be successful, the hearer needs some independent way of identifying the object the speaker is referring to. She needn't be able to locate the object in space, and she needn't even know what *kind* of thing it is. ('Did you hear that?' 'Yes, what on earth was it?') But she needs to have some way of picking out the object that distinguishes it from others, and that doesn't assume that we already know what object the speaker is referring to. (In the example just given, both speaker and hearer might identify the referent of 'that' as 'the distinctive low-pitched wailing sound.' Or they might identify it via perceptually grounded demonstrative thoughts.)

The point is not specific to singular terms. Consider the adverbial demonstrative 'thus.' If I say

(4) While Susan was thus engaged, Bill called the police

I need to ensure that my hearers have sufficient cues to associate 'thus' with some determinate manner of being engaged. These cues may come from the prior discourse context, or from some kind of demonstration, but they must be present. Understanding an utterance of (4) requires having some independent way of saying what is meant by 'thus.' If all you can say is, 'By "thus engaged" he means either *engaged in tying up the burglar* or *engaged in cleaning up the room*, but he definitely doesn't mean *engaged in smoking a pipe*,' then you simply haven't understood my utterance. (At best, you have partially understood it.)

We can ground these observations in more general reflections about the point of truth-conditional semantics. Why do definitions of truth at a context play a central role in orthodox theories of meaning? Because we can explain central aspects of communication by appealing to speakers' and hearers' common knowledge of the conditions for sentences to be true at a context. A speaker who wants to communicate that p chooses a sentence S that is true in the present context just in case p; the hearer trusts the speaker to be telling the truth, and infers that since S is true in the context, p must be the case (Lewis 1980, sec. 2). This picture assumes that speaker and hearer have shared knowledge of what it takes for the sentence to be true in the present context.[1] When that assumption breaks down, truth conditions lose their explanatory relevance.[2]

Let us now apply our point to gradable adjectives. If Kennedy is

---

[1]Or, in cases of 'accommodation' (Lewis 1979), in a slightly adjusted context.

[2]The points made by Stalnaker (1978) complicate this picture somewhat, but not, I think, in a way that spoils the point I am making here. Stalnaker observes that, in order to conform to the pragmatic constraint that it should be common knowledge which proposition has been asserted, participants in a conversation will sometimes take the speaker to have asserted the 'diagonal proposition' expressed by the sentence used. Among the cases that call for diagonalization are cases in which it is not common knowledge what a singular term refers to (Stalnaker's example is 'That is either Zsa-Zsa Gabor or Elizabeth Anscombe'). But in these cases, there must be common knowledge which object the speaker is referring to in each world (considered as actual), or it will not be common knowledge which proposition is the diagonal. In practice this requires that the speaker and hearer have the object in mind under some shared mode of presentation (for example, as 'the woman laughing loudly in the next room'). I will be arguing that this constraint is not met for the contextual delineation functions in Kennedy's semantics.

right that the truth at a context of sentences containing gradable adjectives is sensitive to a contextually determined s function, then knowledge of what it takes for such sentences to be true will require shared knowledge about which function s is. However, as we have seen, Kennedy thinks that we are ignorant of 'the precise location of $s(g)$ in the context.' We do not know, for example, whether s picks out 87mm, 88mm, or 89mm as the minimum diameter for large apples. This means that on Kennedy's view, speakers do not know what it takes for sentences like

   (5) Apple A is large

to be true in a particular context. *A fortiori*, they cannot satisfy the requirement that comes with the use of flexible words—that of giving adequate cues to hearers to allow them to coordinate on the same truth conditions. They are like speakers who say, in front of a mess of colored pills, 'This one is red,' but when asked which pill they mean to refer to can only say, 'one of those red ones, I'm not really sure.' Since truth conditions are only explanatory to the extent that they are known (and known in common) by speakers and hearers, Kennedy's invocation of the s function cannot be doing any explanatory work. It gives us truth at a context, but not in a way that is useful for explaining our use of sentences containing gradable adjectives.

Kennedy might reasonably object that one can know which function s is without knowing at all whether $s(\text{LargenessOf}\,|_{\text{apple}})$ is 87mm, 88mm, or 89mm, or somewhere in between. In general, there are many ways of picking out a function that don't settle all the details of its graph. For example, I might identify a function as 'the Riemann zeta function,' without even knowing what its domain and range are. And I might identify a function as 'the function that maps every person to his or her mother' without being able to give any description of the object to which it maps you other than 'your mother—whoever that may be.'

But what resources do speakers and hearers actually have for picking out the s function, other than by specifying its graph? In the case of the Riemann zeta function, we can simply use the name 'Riemann zeta function,' which has an established conventional meaning. Or we can defer to experts: 'the function mathematicians call the

Riemann zeta function.' But nothing like that is available here: as far as I know, mathematicians have not bothered to give special names to functions that map partial functions from objects to abstract degrees to degrees. In the 'mother of' case, we can specify a rule for determining the value of the function on each possible input. But no such rule presents itself in the case of the *s* function. Kennedy says that *s* 'chooses a standard of comparison in such a way as to ensure that the objects that the positive form is true of "stand out" in the context of utterance, relative to the kind of measurement that the adjective encodes' (2007, p. 17, quoted above). But this is just a *constraint*; it does not come close to *determining s*. Even if we stipulate that, say, only Apple A 'stands out' in the relevant context, there are indefinitely many places where we might draw a line separating A from B.

Delia Graff Fara suggests that we might pick out *s* as the function that maps the semantic value of a gradable adjective to the minimal or maximal degree that is 'significantly greater... than is typical' (Fara 2000, p. 64). As she acknowledges, this way of defining the function is context-sensitive in two ways: '*Significantly greater than* is a context-dependent relation, since what is significant to one person may not be significant to another' (Fara 2000, p. 75). And 'typical' is context-dependent because it can be used in reference to many different kinds of 'norm' (what is expected, what is common, what is optimal, etc.). So the suggestion would be that speakers, in using 'large,' must have in mind (a) a particular person or persons to fill in the *x* in *significant to x* and (b) a particular 'norm' or notion of typicality. I find this implausible, especially since we will need to recognize many, often only slightly different 'norms' of size in order to accommodate the flexibility of 'large.' If Fara's account is correct, then it should be fair game to ask the speaker who uses 'large apple' whether she meant *significantly bigger than the average apple*, or *significantly bigger than you would expect an apple to be*, or *significantly bigger than a healthy apple*, and so on. It would also be fair to ask her, *significantly bigger for whose purposes?* Though she might be able to answer these questions on some reflection, I see no reason to believe that she had to have definite answers in mind in order to make her assertion. And even if she did, she certainly couldn't expect her hearers to divine these answers. A

further problem is that 'significantly' is itself gradable: speaker and hearer might agree about what differences are more significant than others, without agreeing about where to draw the cut-off point for being significant *tout court*.

Hence I can see no way of identifying a particular *s* function that speaker and hearer can plausibly coordinate on. Of course, we can say that it is 'the function that maps the semantic value of a gradable adjective onto the minimum degree on the scale that is required for being characterized with the positive form.' (Or, less technically and more concretely, 'the function that fixes the minimum size something must be to count as "large" in this context.') But this isn't an independent way of picking out a function, any more than 'the thing I'm referring to by "this" ' is an independent way of specifying the referent of 'this.'

It might be thought that the problem is Kennedy's embrace of Williamson's idea that, once contextual factors are settled, 'large apple' has a classical extension, with a sharp cutoff point between the large and non-large apples. Instead of saying that the *s* function yields a definite cut-off point, we might try saying that it yields a 'fuzzy' cut-off point, mapping degrees of largeness to degrees of truth between 0 and 1, with 1 corresponding to the definitely large cases and 0 to the definitely non-large ones. But this doesn't help; indeed, it makes the problem worse. If it is implausible that speakers and hearers coordinate on a particular sharp delineation for 'large,' it is even more implausible that they coordinate on a particular fuzzy delineation. Non-fuzzy properties just need to settle where the dividing line lies. Fuzzy properties need to settle much more: the precise shape of the graph mapping degrees of largeness to degrees of truth. So coordinating on a fuzzy delineation is going to be even harder than coordinating on a sharp delineation.

Here is the upshot. While in using a bare demonstrative like 'this,' one must have a definite object in mind, and successful uptake requires recognizing what object that is, there are no analogous requirements for the use of 'large.' The speaker need not have in mind a particular delineation (even a 'fuzzy' one), and the hearer need not associate the speaker's use with a particular delineation. What we get instead are *constraints* on delineations. In saying that Apple C is large, I rule out certain ways of drawing a line between

large and non-large apples, while leaving others open.

If this is right, then we should not look for a theory of meaning that defines truth conditions (even fuzzy ones) for sentences in context. After all, different ways of drawing the line between large and non-large (different *delineations*) will yield different truth conditions. So if all the speaker has done is *constrain* these ways, without determining one, we don't have enough for truth conditions.

Further reason to think we are only constraining delineations, rather than determining them, is that we regard our options for continuing the conversation as very much open to future decision. If I say 'Apple A is large,' I may leave it open that the conversation may continue in either of these ways:

(6) Yes, but Apple D isn't.

(7) Yes, and Apple D is too.

If asserting that Apple A is large required picking out in thought a particular size property (perhaps a fuzzy one), then at most one of these continuations could be correct.[3]

On the view I am suggesting, our relation to delineations is like our relation to possible worlds, not to propositions. In asserting a proposition we don't commit ourselves to particular world's being actual, but we constrain the worlds we regard as open possibilities. Similarly, I want to say, in asserting a proposition we don't commit ourselves to a particular delineation (even a fuzzy one), but we constrain the delineations we regard as open *practical* possibilities.

It is common to attribute vagueness to 'semantic indecision.'[4] The view I am recommending takes this idea literally and follows it

---

[3]I have used 'Yes' and VP ellipsis to rule out contextual shifts in the property expressed by 'large'; see Stanley (2003). Kennedy seems to think that this problem is averted if *s* is not attached to a particular free variable in logical form (Kennedy 2007, p. 17 n. 15), but it is not clear why this should be. Here the points about agreement and disagreement that I make in MacFarlane (2014 ch. 6) may be relevant.

[4]The locus classicus (but not the origin of the phrase) is Lewis (1986, p. 213): 'The reason it's vague where the outback begins is not that there's this thing, the outback, with imprecise borders; rather there are many things, with different borders, and nobody has been fool enough to try to enforce a choice of one of them as the official referent of the word "outback." Vagueness is semantic indecision.'

through to its natural conclusion. Indecision is a *practical* state; it concerns plans and intentions, not belief. Just as one might plan to buy toothpaste, but not yet have settled on which toothpaste one will choose when confronted with a rack of them at the store, so one might have settled on counting apples greater than 84mm in diameter as large, without having settled on whether one would count a 78mm apple as large. So the mental state of judging an apple to be large is not purely a state of belief; it is, at least partly, a planning state. Gibbard's framework is thus a natural setting for understanding such states and the language we use to express them.

<p style="text-align:center">IV</p>

Planning states and the common ground. A planning state is an action-guiding state of mind with world-to-mind direction of fit— a bit like an intention, but a highly conditional one. The content of a planning state is a *plan.* A plan, as Gibbard conceives it, is something that rules out or permits courses of actions in a possible circumstance. A fully determinate plan, or *hyperplan,* would do this exhaustively for every course of action in every conceivable circumstance. But ordinary planning states fall far short of this. I might plan to go to Berlin, without having decided whether to take the train or a bus. Even if I have decided exactly which bus to take, I probably haven't decided whether to step onto the bus with my left foot first or my right foot. And I certainly haven't decided what to do if I wake up one day and find that I have magical powers.

Both belief states and planning states, Gibbard thinks, stand in relations of compatibility and incompatibility. Belief states are incompatible when they settle certain factual questions differently, and planning states are incompatible when they recommend different actions in the same contingency. Importantly, though, there is no incompatibility between a belief state B and another more settled belief state B^+ that agrees with B on all the questions B settles, but also settles some questions that B left open. Similarly, there is no incompatibility between a planning state P and a more settled planning state P^+ that agrees with P about what to do in all the contingencies for which P recommends a course of action, but also gives recommendations for some contingencies about which P is silent. Firming up one's plan—deciding that one will take the bus,

for example, when one previously left this open—does not count as a *change of mind*, for Gibbard: the later state does not *disagree* with the earlier one.

Given this way of understanding plans, we can represent the content of a planning state as a set of hyperplans (the fully determinate plans that are compatible with it), just as we can represent a belief state as a set of possible worlds (the fully determinate belief contents that are compatible with it). But many mental states are neither purely belief states nor purely planning states: for example, disjunctive states like

(8)   Either apples of diameter 84mm are large or the average size of apples in this plant is greater than 80mm.

To accommodate such states, Gibbard concludes, we need to represent the contents of mental states as sets of world/hyperplan pairs. The pair (w, h) belongs to the content of a state S if the fully decided state that settles on w and h would not disagree with S. This gives Gibbard a solution to the problem of mixed states posed by Geach (1960). We understand such states by seeing which combinations of pure states they rule out (disagree with) or agree with. And we can define basic logical notions just as in standard possible-worlds semantics: a content p entails a content q iff every world/hyperplan pair in p is also in q; p is incompatible with q iff there is no world/hyperplan pair that belongs to both.

This basic Gibbardian idea can be integrated with the style of pragmatics pioneered by Robert Stalnaker (1978, 1999). Stalnaker's idea is that conversations are governed by a 'common ground,' which we can think of either as a set of accepted propositions or as a set of worlds that are regarded as open possibilities. Assertions can be thought of as proposals to add their contents to the common ground. If as assertion that p is accepted, the common ground is modified by removing all worlds at which p is false, so that p now counts as one of the shared assumptions governing the conversation.

This picture assumes that the contents of assertions can be modeled as sets of worlds. If Gibbard is right, though, contents must be modeled as sets of world/hyperplan pairs. It is natural, then, to think of the common ground as a set of world/hyperplan pairs, embodying both shared assumptions about the facts and shared plans for how

to act in various situations. As conversation progresses, the common ground gets more determinate along both the doxastic and the planning dimension.[5]

<div align="center">V</div>

An expressivist semantics. With this framework in the background, let us ask how Kennedy's semantics for gradable adjectives might be modified to capture the idea that asserting 'Apple A is large' can function as a constraint on plans.

As presented above, Kennedy's semantics gives the following truth conditions for the comparative 'larger' and the positive form 'large':

$$[\![\text{larger}]\!]^c = \lambda y.\lambda x.\text{LargenessOf}(x) > \text{LargenessOf}(y)$$
$$[\![\text{large}]\!]^c = \lambda x.\text{LargenessOf}(x) \geq s_c(\text{LargenessOf})$$

If we add a world of evaluation and make the LargenessOf function a function from worlds to degree functions (since an object may be a different size in different worlds), we get something like

$$[\![\text{larger}]\!]^{c,w} = \lambda y.\lambda x.\text{LargenessOf}(w)(x) > \text{LargenessOf}(w)(y)$$
$$[\![\text{large}]\!]^{c,w} = \lambda x.\text{LargenessOf}(w)(x) \geq s_c(\text{LargenessOf}(w))$$

The problem we identified in §III was that speakers and hearers have no way of singling out s_c in thought. The proposal we sketched there was to give up the idea that assertions presuppose a contextually determined delineation, thinking of them instead as *constraining*

---

[5]I am indebted to Yalcin (2012) for showing how easily an expressivist can adapt Stalnakerian pragmatics. However, my implementation differs from his in one important respect. Yalcin proposes that we think of the common ground as a pair consisting of an information state (a set of worlds) and a set of hyperplans. He understands a deontic assertion as a constraint of the second component of this pair. But then how do mixed assertions like (8) or 'Either Moriarty is not near or I ought to pack now' constrain the common ground? These constrain *combinations* of beliefs and plans, so their effect on the common ground cannot be modeled as a constraint on a set of open worlds, a set of open hyperplans, or both. So, instead of factoring the common ground into two pieces, a set of worlds and a set of hyperplans, I propose we think of it in just the same way as we model states of mind: as a set of world/hyperplan pairs.

plans for a delineation. Here is how that might look in our Gibbardian framework:

$$\llbracket \text{larger} \rrbracket^{c,w,h} = \lambda y.\lambda x.\, \text{LargenessOf}(w)(x) > \text{LargenessOf}(w)(y)$$

$$\llbracket \text{large} \rrbracket^{c,w,h} = \lambda x.\, \text{LargenessOf}(w)(x) > s_h(\text{LargenessOf}(w))$$

where $s_h(g)$ = the minimum degree that h recommends counting as satisfying the positive form of g. This clause adds h (for hyperplan) to our points of evaluation. The LargenessOf function, however, is still independent of h. The relative ranking of items on the largeness scale is independent of plans. The only place where h plays a role is in the determination of the cutoff point. Where in Kennedy's semantics s is determined by the context, here it is determined by the hyperplan. Any fully determinate contingency plan will contain plans for where exactly to draw the line between 'large' apples and others.[6]

Since a speaker's state of mind is generally compatible with many hyperplans, this semantics allows that nothing about the speaker's state of mind settles which delineation to use. Accordingly, it does not issue in truth values for sentences at contexts. While a theorist like Kennedy can say

$$\phi \text{ is true at a context } c \text{ iff } \llbracket \phi \rrbracket^{c,w_c} = 1$$

where w_c is the world of the context, we cannot say

---

[6]A complication: Gibbardian hyperplans map subjectively individuated circumstances to sets of permissible actions, and I have not said *at which circumstance h* recommends counting objects with a degree of largeness $\geq$ $s_h(\text{LargenessOf})$ as satisfying the positive form. Nor is it obvious how to do this, with the ingredients represented in the semantics. For purposes of this paper, we can suppose that for each world/hyperplan pair (w, h) in the common ground, h recommends a cutoff point unconditionally, in a way that does not depend on the circumstances. A more adequate resolution of this problem might be to use *delineations* in place of hyperplans in the semantics. We would then have to explain what "delineating states" are without piggybacking on a prior understanding of planning states, but most of the central claims in the paper would survive this change. I am grateful to Seth Yalcin and Robbie Williams for useful discussion here, and I hope to resolve this issue in a more satisfactory way in future work.

$$\phi \text{ is true at a context } c \text{ iff } [\![\phi]\!]^{c,w_c,h_c} = 1$$

because there is not, in general, a 'hyperplan of the context.'[7] Nor can we associate sentences in context with standard possible-worlds intensions. But we *can* associate sentences in context with Gibbardian intensions—sets of world/hyperplan pairs:

$$\text{The intension of } \phi \text{ at } c \text{ is } \{(w,h) \ : \ [\![\phi]\!]^{c,w,h} = 1\}$$

And this is arguably enough to make our semantic theory part of a theory of communication. It tells us what states of mind speakers are expressing in making assertions using sentences, and how the common ground will change if their assertions are accepted.

VI

The dynamics of vagueness. To get a feel for the way vague assertions change the common ground, imagine that when your supervisor begins your training at the apple sorting plant, there are eleven apples are on the table, with their diameters in millimeters marked below them, as in Fig. 1. We can assume it's commonly understood that by 'large' the supervisor means 'large for an apple in this plant', or something similar. Still, at the beginning of your training you don't know much about the kinds of apples they have at this plant. So at the beginning of the conversation, you don't have a very determinate plan for drawing the line between large and non-large apples. As far as you're concerned, it's not ruled out that all of the apples on the table count as medium (that is, neither large nor small). Perhaps there are other apples in the plant (maybe just a few) that are significantly larger and significantly smaller, and it wouldn't be unreasonable to reserve 'large' and 'small' for those. On the other hand, it also wouldn't be unreasonable to use 'large' for the apples above 80mm, and small for those 62mm and smaller. There are some largish gaps there that look like they might be natural joints. As far as you're concerned, there are many reasonable decisions one might make about what to count as a large apple—and your

---

[7]Compare Yalcin (2011), 329, on information states.

supervisor knows this. So the common ground does not initially constrain delineations much at all.

When the supervisor picks up Apple C and says, 'This is a large one,' you regard this as a proposal to plan to count apples that size and larger as 'large.' Since she is the supervisor, and you have an interest in converging with her usage, you accept the assertion, and the common ground is adjusted accordingly. The common ground is now incompatible with plans that fail to count apples with diameter 84mm or greater as large. All of the world/hyperplan pairs in the common ground will now count Apples A–C as large. Using Stalnaker's terminology, we may say that it is *accepted* that Apples A, B, and C are large. However, the claim that Apple F is large is neither accepted nor rejected; the common ground is compatible with delineations that count Apple F as large and with delineations that do not.

Since we already know that Apple C is 84mm across, adding the claim that Apple C is large to the common ground doesn't reduce uncertainty about the apple's size; it serves solely to constrain plans about what is to count as large. By contrast, 'Some apples are less than 34mm across' would constrain only worlds, not plans. But many assertions will constrain both worlds and plans in an intertwined way. For example, suppose that the supervisor asserts

(9) Most of the apples in this plant are large.

Accepting (9) may not rule not out any particular worlds or plans. But it does rule out many *combinations* of worlds and plans: for example, the combination of a world where most of the apples are less than 80mm diameter and a plan to count apples as large only if they are more than 80mm in diameter.

On the picture I have been sketching, delineations are constrained by our assertions in a way that is *transparent*. When the supervisor says that Apple C is large, and this is accepted, the effect on the common ground is clear: we throw out all *and only* hyperplans that fail to count Apple C (and larger apples) as 'large.' So, a hyperplan that fails to count Apple D as 'large' is compatible with the adjusted common ground, even though Apple D is only 1mm less in diameter than Apple C. This result may seem controversial or suprising, and I will defend it in the next section. But it should be clear why I need

to accept it. If the common ground is to be truly 'common,' the effect of assertions on the common ground cannot be uncertain. We cannot say (as an epistemicist might) that calling C large rules out delineations that fail to count apples $(84 - \delta)$mm in diameter as large, for some unknown δ. Nor can we say that the affect on the common ground is indeterminate. Either of these moves would bring back the problem we found in Kennedy's view: we would be unable to explain how speakers and hearers coordinate on a determinate update to the common ground.

<div align="center">VII</div>

Vagueness as a pragmatic phenomenon. To say that the supervisor's assertion that Apple C is large leaves open delineations that fail to count Apple D (only 1mm smaller) as large is to allow that she might go on to assert that Apple D is not large. Were this assertion accepted, the common ground would put a strong constraint on delineations, requiring a boundary between medium and large apples to be drawn somewhere in the space between 83mm and 84mm. This seems to go against the idea that 'large' and other vague words are distinctively 'tolerant': that tiny differences in size cannot make a difference between being a large apple and not being one.[8]

Tolerance certainly is, as they say these days, 'a thing.' But what kind of thing? We know that, on pain of paradox, it cannot be a *semantic* entailment that an apple 1mm smaller than a large apple is large. Tolerance must, then, be explained pragmatically, as a feature of our *use* of vague words. What we need to explain is why it normally seems wrong, having called one apple large, to class another apple just 1mm smaller as not large.

On the present view, this becomes a question about the reasonableness of proposing to constrain our future uses of 'large' in a way that removes nearly all of our flexibility to refine what counts as 'large' to fit our developing needs. A proposal like this might be reasonable in certain cases: when it's important that every apple be classified as small, medium, or large, with no undecided cases, and that this

---

[8]The term is due to Crispin Wright, who defines it thus: 'F is *tolerant* with respect to ϕ if there is also some positive degree of change in respect of ϕ insufficient ever to affect the justice with which F applies to a particular case' (Wright 1976, p. 229).

be done consistently; when it is clear from the beginning what the most useful criteria for this classification will be; when these are criteria that the people doing the sorting can use effectively; when it is not important that the classification cut at 'natural joints' that might only be discerned with experience; when there are no fixed limits on the number of apples that can fit in each category. These conditions might actually be met in the apple sorting plant, and in this context, I don't think it would be outrageous for the supervisor to declare that Apple C is large and Apple D medium. This would be a way of telling you exactly how big an apple needs to be to count as large (compare Tappenden 1993, pp. 565–6).

However, most ordinary situations are not like this. Ordinarily, we have reason to retain a certain flexibility in how we draw our boundaries, so that we can be guided in doing so by future experience and evolving needs. Imagine, for example, that we are working in the admissions office of a college. We want to classify some students as *In* and the rest as *Out*. Ultimately we'll have a dichotomy, with every student in one of these categories, but we'll start by designating some clear cases as *In* and *Out*. As we move away from the clearest cases, we need to retain some flexibility, and it would be contrary to our purposes at this stage to propose some particular combination of grades and SAT scores as the cutoff point for *In* status. We might prefer to admit between 90 and 100 students, for example, unless the pool is really outstanding—and a cutoff point chosen without looking at all of the files could lead us to admit more. In addition, where we draw the line may depend on the composition of the pool. If there are a lot of applicants clumped around one score, and the next best applicants score far lower, we might find that a reason to draw the line below the clump. So if, early in the process, someone proposes a particular sharp cutoff line, the proposal to constrain delineations in this way will be rejected for practical reasons.

Jamie Tappenden reminds us of cases where legislators or judges intentionally avoid premature line-drawing in order to allow future application of the law to be informed by an evolving appreciation of the situation (Tappenden 1994, p. 198). In other cases, our pragmatic reason for refraining from drawing lines is rooted in our perceptual limitations. In a context demanding visual classification of colors, a proposal to draw a sharp line between orange and red would

be rejected as a bad idea because we lack the ability to locate this line reliably on a spectrum. But in a context where we are able to measure wavelengths with instruments, or in which cans of paint are marked with numbers, such a proposal might be reasonable—depending on our purposes.

I suggest, then, the 'tolerance' intuition can be explained as an awareness that a proposal to draw a sharp line in any particular place would be rejected for pragmatic reasons. Nothing about the *meanings* of vague words is inconsistent with drawing a sharp boundary; it's just that the cases in which a proposal to draw a sharp boundary would be sensible are few and unusual.

Since it is undeniable, I think, that it might be perfectly sensible for your supervisor in the apple sorting plant to declare that Apple C is 'large' and Apple D is not, and that we would know how to go on using language in a way consistent with this, anyone who thinks that this usage is ruled out by the *meaning* of 'large' must say that the supervisor has used 'large' in a special, technical sense that diverges from its ordinary sense. This response is explicit in Fodor and Lepore (1996), 525–6, who class these cases with

(10) Ketchup counts as a vegetable for the purposes of the Republicans' school-lunch program.

But once we accept that we can make the boundary between 'large' and 'medium' more determinate by making classifications that are accepted, it is very hard to draw a line between precisifications that change the meaning and those that do not. And, while we could follow (10) by saying 'but of course it isn't a vegetable,' it would be very odd to say that although Apple D counts as medium for sorting purposes, it 'isn't really medium.' I am in agreement with Fine (1975, p. 277), then, that 'language can retain its identity upon precisification.'

So far I have been focused on why we *resist* drawing a boundary between 83mm and 84mm. One might also ask why, having accepted that an apple 84mm across is large, we tend to accept the subsequent claim that an apple 83mm across is large. ('Forced-march sorites arguments' are just sequences of such assertions, each of which we feel we must accept given our acceptance of the last.) Presumably, the reason is this: having acknowledged by our acceptance

that the earlier proposal was reasonable, it is hard to see a reason why the new, only slightly different one should be rejected as unreasonable. To say it's unreasonable, one would need to have a clear purpose in mind for which that 1mm difference really matters. It is hard enough to find such purposes in ordinary situations, and even harder in the artifical forced march sorites situations philosophers think about, where there is no clear purpose to the classifications.

<div align="center">VIII</div>

The truth in nihilism. Nihilism is the view that sentences with vague terms are incapable of being true or false. The view is usually ascribed to Gottlob Frege, with some justice. Frege held that the *Bedeutung* (reference) of a sentence is its truth value (Frege 1979, pp. 122, 194–5, 1984, pp. 144–7, 162–5), that the *Bedeutung* of a sentence is a function of the *Bedeutung* of its parts (Frege 1979, pp. 192–4, 1984, pp. 162–3), that a predicate only has a *Bedeutung* if it sharply divides the objects into the universe into two classes, those that fall under it and those that do not (Frege 1979, p. 122), and that vague predicates do not do this (Frege 1979, p. 155, 1980, p. 114). It follows directly that vague sentences cannot be true or false. In a letter to Peano, Frege concedes that in ordinary life, we do manage to use vague sentences to express thoughts capable of truth or falsity (Frege 1980, p. 115; Puryear 2013, p. 133). But his account of language leaves us with no explanation of how we do that, and no criterion of validity for inferences involving vague expressions.

The expressivist view vindicates something close to Frege's view that vague predicates lack extensions. On this view, predicates have extension relative to worlds and hyperplans, and the context is almost always compatible with hyperplans that make different delineations. So, even given full information about the context and the factual state of the world, there is simply no answer to the question, 'what is the extension of "large for an apple" '? Similarly, until we are given both a world and a hyperplan, there is no answering the question whether 'That apple is large' is true. This view is 'non-factualist' about vague claims in just the same way that Gibbard's view is non-factualist about normative claims.

We should not say, however, that vague predicates lack *Bedeutung*. If we understand the *Bedeutung* of an expression as its compositional

semantic value, then vague predicates have *Bedeutungen.* Unlike Frege's, our theory gives a perfectly clear account of the semantic contribution vague predicates make to sentences, and of the inferential relations such sentences stand in.

It might seem that our theory has unacceptably revisionary consequences. Do we not apply 'true' and 'false' to vague sentences all the time?

(11) It is true that bald men are prone to getting sunburns on their heads.

But, like Gibbard, we can make good sense of an ordinary, monadic truth predicate that obeys the propositional T-schema:

(12) The proposition that p is true iff p.

The semantics is simple: the extension of 'true', at a world/hyperplan pair (w, h), is the set of propositions that hold at (w, h).

As a rehabilitation of nihilism, expressivism has some advantages over the account of Braun and Sider (2007). Braun and Sider deny that sentences containing vague terms express unique propositions. Rather, they are associated with a cloud of propositions expressed by the 'legitimate disambiguations' of the sentence used. When we are 'ignoring' a sentence's vagueness, Braun and Sider say, we are justified in using it to make an assertion when it is 'approximately true'—meaning that all of its legitimate disambiguations are true. But why should truth on all disambiguations amount to approximate truth, and why should it be a standard for correct assertion? Suppose I say, 'Bonny went to the bank,' ignoring the ambiguity of 'bank,' and not intending to use it in either of its two senses. And suppose Bonny went to the money bank *and* the river bank, so that my sentence is true on all disambiguations. Isn't my utterance simply a failure as an assertion? On the expressivist account, by contrast, we use vague sentences to assert vague propositions, which have distinctive effects on the common ground, and which can be believed, doubted, or denied just like other propositions. Braun and Sider's mistake is much like Frege's: they think a compositional account of language must assign factualist truth conditions to sentences.

The nihilist's insight—that factualist truth conditions are not going to give us an adequate compositional account of the meanings of

vague sentences—is perfectly correct. Where the nihilist goes wrong is in thinking that factualist truth conditions exhaust our resources for giving such an account.

<div align="center">IX</div>

The truth in supervaluationism. According to supervaluationism, vague sentences are true if they are true on all acceptable classical precisifications that respect clear cases and analytic truths ('penumbral connections'). The expressivist view has a similar upshot: vague sentences (or propositions) are *accepted* if they hold in every world/hyperplan pair in the common ground. Since a hyperplan will recommend a cutoff point for the application of the positive form of every vague predicate, this means that vague propositions are accepted if they hold on every classical precisification that satisfies the constraints imposed so far in the conversation. Thus it may look as if the expressivist view is just supervaluationism by another name. The differences, though, are important, and they favor the expressivist approach.

On the supervaluationist view, a disjunction can be true even though neither disjunct is, and an existentially quantified sentence can be true even though no instance is. These anomalies have often been taken to be problematic for supervaluationism. The analogous expressivist claims, about acceptance rather than truth, are innocuous. A disjunction can certainly be accepted though neither disjunct is: that is what happens when we know that Smith or Jones committed the murder, but we don't know which. Similarly, an existentially quantified sentence can be accepted even when no instance is. It can be accepted that someone committed the murder, even though there is no person such that it is accepted that that person committed the murder.

The difference between truth and acceptance also becomes important when we consider the status of 'in between' sentences—those that are neither true nor false, or neither accepted nor rejected. On the supervaluationist view, a borderline sentence is not true, and so presumably ought not be asserted, for the same reason that other non-true sentences ought not be asserted. So, for supervaluationists, the borderline area is a forbidden zone. On the expressivist view, by contrast, the borderline area is an area for potential exploration.

There is nothing wrong with asserting propositions that are not yet accepted; indeed, in the normal case one asserts propositions that are not yet accepted, with the aim of making it the case that they are accepted.

Supervaluationism is right, then, to see that acceptance in a common ground depends on truth throughout a range of acceptable delineations, but it distorts this insight by identifying this property with truth.

<div align="center">X</div>

The truth in epistemicism. According to epistemicists, 'large (for an apple)' has a classical extension at a context; it exhaustively partitions apples into two classes, the large and the not-large. Where the boundary lies, we don't know—and can't know, because of 'margin of error' requirements on knowledge. The distinctive phenomena of vagueness can be explained in terms of this inescapable ignorance. Thus understood, vagueness puts no pressure on classical semantics or logic.

The obvious objection to epistemicism is that it requires a completely opaque relation between meaning and use. Why does 'large' draw a line where it does, and not somewhere else? Not because we intend it to draw a line there—since we don't even know where the line is. Somehow, Williamson insists, the facts about our use of 'large' must determine where this line lies. But the way it does so is inscrutable to us. This is hard to swallow.

However, epistemicism has a powerful argument in its favor: *every* plausible view of vagueness seems to be committed to inscrutable semantic boundaries. For example, the supervaluationist will have a boundary between the smallest apple of which 'large' is true and the biggest one of which 'large' is neither true nor false. The degree theorist will have a boundary between the smallest apple of which 'large' is true to degree 1 and the biggest one of which 'large' is true to degree less than 1. And so on. So, why bother with a nonstandard semantics and logic if one is going to end up with inscrutable semantic boundaries, and an opaque relation between meaning and use, anyway?

The expressivist view acknowledges that there will always be hid-

den semantic boundaries. We don't know, to the micron, where the line between large and medium apples lies. But we must distinguish two kinds of unknown boundaries. One kind, which the epistemicist is committed to, is problematic because it makes it impossible to understand *how* our use of words in context determines the boundaries. The other kind—the only one the expressivist view requires—is innocuous.

Ignorance about semantic boundaries is innocuous when it results from our intentionally anchoring these boundaries to particular worldly facts of which we are ignorant. Your supervisor holds up an apple and says, 'This is a small one.' Assuming the assertion is accepted, your supervisor has established a semantic boundary: it is settled that any apple smaller than that apple counts as 'small' (bigger ones may still be up for determination). But you don't know how big the apple is, exactly: it might be anywhere between 55mm and 65mm and diameter. So you don't know whether this newly established semantic boundary is at 58mm or 60mm or 62mm.

This is a kind of ignorance that will affect most uses of vague words. But it does not raise any mysteries, I take it, about how the semantic boundaries arise from our usage. It is the same kind of ignorance early chemists had when they thought of water as the natural kind that was the dominant constituent of lakes and rivers, but did not know it was composed of H_2O molecules. Suppose you discover that the apple the instructor calls 'small' is exactly 58mm in diameter. Then, on the expressivist account, you know exactly where the line is between the settled cases and the unsettled cases: it is at 58mm.

The problematic unknown boundaries—call them *inscrutable* boundaries—are those that cannot be explained by appeal to our meaning intentions and our ignorance of wordly facts to which these intentions make reference. These boundaries would remain unknown even if we knew all the facts about speakers' intentions, about what has been asserted in the conversation, and about the sizes of the relevant apples. Epistemicists like Williamson are committed to inscrutable boundaries. But so are many non-epistemicists: for example, supervaluationists who say that the notion of 'legitimate interpretation' is itself vague (Keefe 2000), three-valued theorists who use a vague metalanguage (Tye 1997), and contextualists who take our uses of vague words to push

the semantic boundaries a little beyond where they need to be for our assertions to be true (Soames 1999). Williamson's insight is that his acceptance of inscrutable boundaries cannot be a dialectical disadvantage against other views that accept such boundaries (Williamson 1994, 2002). Accepting this insight, the expressivist view rejects inscrutable boundaries. It ties semantic boundaries to our usage in a way that is transparent and deterministic, and thinks of ventures into the borderline zone as proposals rather than unjustified guesses.

<div align="center">XI</div>

The truth in contextualism. The expressivist account has much in common with with dynamic and contextualist accounts of vagueness (Raffman 1996, 1996; Soames 1999; Fara 2000; Kyburg and Morreau 2000; Barker 2002; Shapiro 2008). It agrees with them in emphasizing the *flexibility* of vague terms as the key to the phenomena distinctive of vagueness. But although these views get almost everything right, their central insights are distorted by the mandate to deliver truth values in context.

Fara (2000), as we have seen, secures truth values in context by attributing intentions to speakers that outstrip what they seem to have in mind or what hearers are in a position to discern. Kyburg and Morreau (2000), who appreciate the way in which uses of vague words function as constraints on delineations, resort to supervaluations in order to assign truth values to sentences in context (582). And Barker (2002), who sees clearly that sentences like 'Feynman is tall' can have the non-descriptive role of constraining delineations, blurs the insight by treating the delineation as determined by the world or context.[9] This forces him to conceive of assertions like 'Feynman is tall' as reducing our 'ignorance' (3–4) or 'uncertainty' (9) about 'the' delineation governing the discourse, rather than constraining our plans.[10]

---

[9]'d will be a function that maps a world onto the delineation that characterizes the vague predicates in use in the discourse of that world.... $d(c)(\llbracket \text{tall} \rrbracket)$ yields the standard of absolute tallness in c' (Barker 2002, p. 6). Barker's d is essentially the same as Kennedy's s, criticized above.

[10]Interestingly, although Barker's semantics assigns a determinate classical extension to the positive form of 'tall' at a context, $d(c)(\llbracket \text{tall} \rrbracket)$, the only use it

Soames (1999), Raffman (1996), and Shapiro (2008) explain tolerance intuitions by arguing that, although there are always going to be semantic boundaries (either between the extension and the antiextension, or between the extension and an undetermined borderline area), these boundaries are never where we are looking: as we move along a sorites sequence, the extensions of vague terms shift. We confuse our inability to exhibit a boundary with the absence of such boundaries. But it is difficult for such views to explain our sense of common subject matter. As Stanley (2003) points out, we would naturally say, 'This patch is red, and this one *is too*, and this one *is too*, ...' The 'is too' would be inappropriate if there were gradual contextual change in the property expressed by 'red'.

The expressivist view does much better here. It does not need to posit any shift in the property expressed by a vague term; the term can have the same intension throughout a conversation. What changes is the common ground. As we move along the sorites series, certain hyperplans that were previously open get ruled out, and propositions that were not previously accepted become accepted. This process of firming up indeterminate plans is no more a change of mind or shift in subject matter than the process of moving from agnosticism to belief.

To explain tolerance intuitions, the expressivist view need not posit hidden semantic boundaries that move around as we use vague terms. The boundaries can be perfectly manifest. Once we start thinking of vagueness in terms of plans, explaining tolerance is just a matter of explaining why certain plans are almost never going to be accepted. The explanations are practical and require no special story about the meanings of vague words.

Contextualist approaches to vagueness have the right idea, but they put it in the wrong setting. The right setting, I have argued, is a kind of plan-expressivism originally developed for quite different ends. If

---

makes of this is to determine an update function for 'tall' which is essentially a constraint on delineations: an assertion of 'Feynman is tall' rules out every context c such that Feynman's degree of height in c is at least as great as $d(c)(\llbracket\text{tall}\rrbracket)$ (Barker 2002, p. 7). If one thinks of Barker's 'worlds' or 'contexts' as world/hyperplan pairs, then Barker's approach is very similar to the one recommended here. The idea of a contextually determined delineation plays no essential role in his dynamics, and it muddies the philosophical picture.

we think of vague language in this framework, we can capture what is right about contextualist, supervaluationist, nihilist, and epistemicist approaches to vagueness, while avoiding their pitfalls.[11]

REFERENCES

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

Braun, David and Theodore Sider 2007: 'Vague, so Untrue'. *Nous*, 41, pp. 133–156.

DeRose, Keith 2008: 'Gradable Adjectives: A Defense of Pluralism'. *Australasian Journal of Philosophy*, 86, pp. 141–160.

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

Fine, Kit 1975: 'Vagueness, Truth and Logic'. *Synthese*, 30, pp. 265–300.

Fodor, Jerry A and Ernest Lepore 1996: 'What Cannot Be Evaluated Cannot Be Evaluated and It Cannot Be Supervalued Either'. *The Journal of Philosophy*, pp. 516–535.

Frege, Gottlob 1979: *Posthumous Writings*. (Friedrich Kaulbach Hans Hermes, Friedrich Kambartel, Ed.). Chicago: University of Chicago Press.

Frege, Gottlob 1980: *Philosophical and Mathematical Correspondence*. (Brian McGuinness, Ed., Hans Kaal, Trans.). Chicago: University of Chicago Press.

Frege, Gottlob 1984: *Collected Papers on Mathematics, Logic, and Philosophy*. (Brian McGuinness, Ed.). Oxford: Basil Blackwell.

Geach, Peter 1960: 'Ascriptivism'. *Philosophical Review*, 69, pp. 221–225.

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

Keefe, Rosanna 2000: *Theories of Vagueness*. Cambridge: Cambridge University Press.

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

Kyburg, Alice and Michael Morreau 2000: 'Fitting Words: Vague Lan-

guage in Context'. *Linguistics and Philosophy*, 23, pp. 577–597.

Lewis, David 1979: 'Scorekeeping in a Language Game'. *Journal of Philosophical Logic*, 8, pp. 339–359.

Lewis, David 1980: 'Index, Context, and Content'. In Stig Kanger and Sven Öhman (Eds.), *Philosophy and grammar* (pp. 79–100). Dordrecht: Reidel.

Lewis, David 1986: *On the Plurality of Worlds*. Oxford: Oxford University Press.

MacFarlane, John 2014: *Assessment Sensitivity: Relative Truth and Its Applications*. Oxford: Oxford University Press.

Puryear, Stephen 2013: 'Frege on Vagueness and Ordinary Language'. *Philosophical Quarterly*, 250, pp. 120–140.

Raffman, Diana 1996: 'Vagueness and Context Relativity'. *Philosophical Studies*, 81, pp. 175–192.

Richard, Mark 2008: *When Truth Gives Out*. Oxford: Oxford University Press.

Shapiro, Stewart 2008: *Vagueness in Context*. Oxford: Clarendon Press.

Soames, Scott 1999: *Understanding Truth*. Oxford: Oxford University Press.

Stalnaker, Robert 1978: 'Assertion'. In P. Cole (Ed.), *Syntax and semantics, vol. 9: Pragmatics*. New York: Academic Press.

Stalnaker, Robert 1999: *Context and Content: Essays on Intentionality in Speech and Thought*. Oxford: Oxford University Press.

Stanley, Jason 2003: 'Context, Interest Relativity and the Sorites'. *Analysis*, 63, pp. 269–80.

Tappenden, Jamie 1993: 'The Liar and Sorites Paradoxes: Toward a Unified Treatment'. *Journal of Philosophy*, 90, pp. 551–577.

Tappenden, Jamie 1994: 'Some Remarks on Vagueness and a Dynamic Conception of Language'. *Southern Journal of Philosophy*, 33, pp. 193–201.

Tye, Michael 1997: 'Sorites Paradoxes and the Semantics of Vagueness'. In Rosanna Keefe and Peter Smith (Eds.), *Vagueness: A reader* (pp. 281–293). Cambridge: MIT Press.

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

Williamson, Timothy 2002: 'Soames on Vagueness'. *Philosophy and Phenomenological Research*, 65, pp. 422–428.

Wright, Crispin 1976: 'Language-Mastery and the Sorites Paradox'.

In Gareth Evans and John McDowell (Eds.), *Truth and meaning* (pp. 223–247). Oxford: Oxford University Press.

Yalcin, Seth 2011: 'Nonfactualism About Epistemic Modality'. In Andy Egan and Brian Weatherson (Eds.), *Epistemic modality* (pp. 295–332). Oxford University Press.

Yalcin, Seth 2012: 'Bayesian Expressivism'. *Proceedings of the Aristotelian Society*, 112, pp. 123–160.