

Forthcoming in *Cuts and Clouds* (ed. Moruzzi, Dietz, Wright), OUP.

Demoting Higher-Order Vagueness

Diana Raffman

Higher-order vagueness is widely thought to be a feature of vague predicates that any adequate theory of vagueness must accommodate. It takes a variety of forms. Perhaps the most familiar is the supposed existence, or at least possibility, of higher-order borderline cases—borderline borderline cases, borderline borderline borderline cases, and so forth. A second form of higher-order vagueness, what I will call ‘prescriptive’ higher-order vagueness, is thought to characterize complex predicates constructed from vague predicates by attaching operators having to do with the predicates’ proper application. For example, the predicates ‘mandates application of “old”’ and ‘can competently be called “old”’ are prescriptively higher-order vague. Higher-order vagueness appears in other guises as well,¹ but these two have been of particular interest to philosophers and will be my target here. I want to expose some misconceptions about them. If I am right, higher-order vagueness is less prevalent, and less important theoretically, than is usually supposed.²

In what follows I am going to assume that vagueness is a semantic feature of natural language. For the most part I won’t discuss epistemic or pragmatic views, and I will say nothing about so-called metaphysical vagueness.

1. *Higher-order borderline cases*

That vague predicates have or could have higher-order borderline cases is largely taken for granted by theorists of vagueness. On the standard view, first-order borderline cases for a vague

predicate ‘ Φ ’ are neither-definitely- Φ -nor-definitely-not- Φ .³ Second-order borderlines (or anyway one set of second-order borderlines) are then neither-definitely-definitely- Φ -nor-definitely-not-definitely- Φ ; third-order borderlines are neither-definitely-definitely-definitely- Φ -nor-definitely-not-definitely-definitely- Φ ; and so on. This hierarchy of ever higher orders of borderline cases is often said to continue *ad infinitum*, thereby constituting, or at least providing for, the blurred boundaries of the predicate ‘ Φ ’.

There are problems, however. For one, sharp cut-offs reappear in the end. Mark Sainsbury explains:

[S]uppose we have a finished account of a [vague] predicate, associating it with some possibly infinite number of boundaries, and some possibly infinite number of sets. Given the aims of the description, we must be able to organize the sets in the following threefold way: one of them is the set supposedly corresponding to the things of which the predicate is absolutely definitely and unimpugnably true, the things to which the predicate’s application is untainted by the shadow of vagueness; one of them is the set supposedly corresponding to the things of which the predicate is absolutely definitely and unimpugnably false, the things to which the predicate’s non-application is untainted by the shadow of vagueness; the union of the remaining sets would supposedly correspond to one or another kind of borderline case. So the old problem re-emerges: no sharp cut-off to the shadow of vagueness is marked in our linguistic practice, so to attribute it to the predicate is to misdescribe it (255).

Sainsbury's reasoning seems to me decisive; and anyway there are simpler and more plausible ways to understand the blurred boundaries of a vague predicate (e.g., in terms of tolerance or soriticality). In addition, if an infinite hierarchy of borderline cases were required for blurred boundaries, then there would be sharp cut-offs in a sorites series.⁴ That can't be right.

My present aim is to articulate some further, mostly intuitive worries about higher-order borderline cases as standardly conceived. I will do this by setting out a series of informal questions and criticisms—I'll call them 'ruminations'—that help to reveal just how problematic the notion is.

Rumination #1. Consider the set containing all possible borderline cases of any order for vague predicate ' Φ ', as in Sainsbury's "finished account". Why aren't all of these items just first-order borderline cases? Don't they all fall within a gap between the extensions of ' Φ ' and ' $\text{not-}\Phi$ '? Alternatively, why aren't these items just more (first-order) borderlines, definitely Φ items, and definitely $\text{not-}\Phi$ items? A competent speaker could classify each of them in one of these ways, couldn't she? In fact I think we have no grasp at all on the idea of an item that doesn't fit into any of these three categories.

Rumination #2. If a hierarchy of borderline cases doesn't make for blurred boundaries, why else believe in them? If there can be borderline cases between Φ and $\text{not-}\Phi$, the thinking goes, then surely there can be (second-order) borderline cases between Φ and borderline Φ ; and then surely there can be (third-order) borderline cases between Φ and the second-order borderlines; and so on.⁵ This line of reasoning sounds plausible, but it overlooks a crucial possibility: *viz.*, that there can be borderline cases between ' Φ ' and ' $\text{not-}\Phi$ ' only insofar as Φ and $\text{not-}\Phi$ are not themselves borderline categories. It may be that only non-borderline categories can have borderline cases. Notice that borderline cases are defined negatively, in terms of an absence or

lack—specifically, a lack of category membership. Borderlines are possible only insofar as the (definite) extensions of vague ‘ Φ ’ and ‘not- Φ ’ are not together exhaustive over the range of values in a relevant sorites series. This is why we say that borderline cases “fall within the gap” between the extensions of ‘ Φ ’ and ‘not- Φ ’. To put the point another way, there is nothing more to being borderline than failing to (definitely) belong either in the category Φ or in the category not- Φ . Consider how we classify items as borderline: presumably we measure them against, or judge their “distance” from, the definite cases of Φ and not- Φ at the endpoints of a sorites series. But if there were definite borderline cases, surely we would classify items as borderline by judging their distance from *those*. (Indeed I think it is misleading to speak of being borderline as a ‘category’. Better to call it, say, a ‘status’.)

My thought then is that ‘borderline Φ ’ may not be the right sort of predicate to have borderline cases of application; it is not sufficiently centered or anchored, one might say. Thus when we talk of definite and borderline borderline cases, we are no longer treating the items in question as defined negatively, as falling within a gap. We are in effect transforming the (first-order) borderline cases into a new, non-borderline category with its own center of gravity—a full-fledged incompatible of Φ and not- Φ . As evidence of this transformation, consider where the putative second-order borderlines are supposed to be located in a sorites series—say, a series of heights progressing from a definitely tall height (e.g., 6'5") to a definitely average height (e.g., 5'9") compared to British men. And suppose that 5'10½" is a borderline case (B1). (See Figure 1.) Then the second-order borderline cases would be located as in Figure 2. However, I predict that if a competent speaker were asked to proceed along the original tall/average series from the definitely tall height 6'5" to the definitely borderline height 5'9", and to classify each height as definitely tall, definitely borderline, or borderline borderline (B2), she would locate any border-

line borderline cases not as in Figure 2, but roughly as in 3. Among other things, she would now classify as borderline borderline some heights that she previously classified as tall. The span of the first-order borderlines would spread out, as it were, pushing everything toward the tall end.




Figure 1




Figure 2




Figure 3




Figure 4

Figure 3 still isn't right, however; for in classifying the heights in this new, shorter sorites series, the speaker would not in fact be classifying them as tall, first-order borderline, and second-order borderline. Instead, she would be classifying them as tall, first-order borderline, and, say, above average, as in Figure 4. She would be transforming what had been first-order borderline cases into a new height category, above average, with its own, new, first-order borderline cases that are neither-definitely-tall-nor-definitely-not-tall (above average). Unsurprisingly, some heights that were tall when 'tall' was opposed to 'average' are borderline or even not-tall

when ‘tall’ is opposed to ‘above average’. (One might say that the standard for being tall as opposed to above average is higher than the standard for being tall as opposed to average.)

I have not done a study to confirm this prediction, but some support comes from related observations. Sainsbury notes that “subjects asked to classify a range of test objects using just ‘young’ and ‘old’ make different assignments to these words from those they make to them when asked to classify using, in addition, ‘middle-aged’” (1997, 259). C. L. Hardin makes a similar claim about hue predicates:

[T]he boundary of red in the broadest sense extends to the immediate neighborhood of unique yellow, and the breadth of that spread we acknowledge by our use of the modifier ‘reddish’. But, in a somewhat narrower sense, the boundary between red and yellow falls at the point at which the perceptual “pull” of yellow is equal to that of red. This point is, of course, orange. But once we introduce orange as a distinct hue category, its boundary with red is at issue, and the extension of ‘red’ must be contracted to make room for the oranges. The natural red-orange boundary would seem to fall at the 75 per cent red, 25 per cent yellow region which was well within the scope we took ‘red’ to have when we were concerned to compare red with yellow (1988, 184).

I expect that an analogous contraction of the extension of ‘ Φ ’ (e.g., ‘tall’) would occur if a speaker attempted to locate second-order borderline cases in a sorites series.

Rumination #3. Borderline cases are supposed to be of indefinite or indeterminate or uncertain status with respect to being Φ . So borderline cases have a status other than being Φ .

Therefore definite borderline cases definitely have a status other than being Φ . But intuitively,

how could it be indefinite whether an item that is *definitely other than Φ* is Φ ? Allowing that ‘ x is not Φ ’ is true on a weak reading of the negation is not an adequate response, in my view. The intuitive question is: how could it be indefinite whether a definitely borderline item is Φ , rather than just plain false?

Rumination #4. The impossibility of higher-order borderline cases seems to follow from two intuitively plausible claims about vague predicates. For all vague predicates ‘ Φ ’ and ‘ Ψ ’:

- (i) If an item is definitely Φ , then failure to classify it as Φ is mistaken or in some way improper or at least legitimately questionable.⁶
- (ii) Failure to classify an item as borderline Ψ cannot be mistaken or in any way improper or even legitimately questionable. (Intuitively, one is never required to classify something as borderline; a judgment of ‘borderline’ is always optional.)

If (i) and (ii) are true, then (iii) follows straightforwardly:

- (iii) Therefore no item can be definitely borderline Ψ .

Given (iii), (iv) appears to follow (or so I will contend):

- (iv) Therefore no item can be borderline borderline Ψ .

Call this the ‘Simple Argument’. It is so simple that it may seem to involve some sleight of hand; so I want to spell out the justification for each step. First, though, I want to acknowledge that one can of course define ‘definitely’, as a technical term, however one wants. But technical control risks estrangement from the ordinary meaning and application of vague words. The Simple Argument, and my ruminations in general, proceed on the assumption that the meaning of the definiteness operator in a theory of vagueness is grounded in the meaning of ‘definitely’ as

used by ordinary speakers when they apply vague predicates. On this assumption, the behavior of the definiteness operator is in some measure constrained by ordinary linguistic intuition.

Understanding ‘definitely’ in this way, let us consider how the premises and reasoning in the Simple Argument can be justified. Premise (i) makes an extremely weak claim about the character of definitely Φ items. If, contrary to (i), definitely Φ items can also permissibly be classified (e.g.) as not- Φ or as borderline Φ , then it is hard to see what definiteness comes to. Perhaps items that can competently be classified as borderline Φ can also competently be classified as Φ and as not- Φ . But the analogous claim is not plausible for ‘ Φ ’ and ‘not- Φ ’: it is not the case that any item that can competently be classified as Φ (not- Φ) can equally competently be classified as borderline Φ or as not- Φ (Φ). Definitely Φ (not- Φ) items appear to carry some sort of requirement that they be so classified, thus making failure to do so mistaken or at least questionable.

Premise (ii), it seems to me, can be found in ordinary linguistic intuition (but for experimental evidence that competent speakers proceeding along a sorites series do not always—indeed often do not—employ the category borderline even when it is explicitly made available, see Lindsey *et al.*, 2008 [in progress].)⁷ Of course (ii) goes hand in hand with the thought that any item that can competently be classified as borderline can also competently be classified as Φ and as not- Φ . If (ii) is accepted, premise (iii) then follows. Premise (iv) is secured from (iii) not merely because ‘definitely Φ ’ and ‘borderline Φ ’ are interdefinable, but because, as I argued above, ‘borderline Φ ’ is defined by ‘definitely Φ ’ in a wholly negative fashion: there is nothing more to being borderline Φ than failing to be either definitely Φ or definitely not- Φ . Hence if definite borderline cases are impossible, so are borderline borderline cases.

Stewart Shapiro and Elia Zardini have pointed out (in conversation) that if definite borderline cases are impossible, then it seems to follow, absurdly, that all first-order borderlines are second-order borderlines. For if nothing can be definitely borderline, then, trivially, first-order borderlines are not definitely borderline. But first-order borderlines are also not-definitely-not-borderline. Therefore first-order borderline cases are not-definitely-borderline and not-definitely-not-borderline, which is just the definition of a second-order borderline case.

The trouble with this clever objection is that if definite borderline cases are impossible, then second-order borderline cases are also *not (first-order) borderline*. Consider: first-order borderline cases come between the definitely Φ items and the definitely not- Φ items, being neither one nor the other ($\sim\text{Def}\Phi x$ & $\sim\text{Def}\sim\Phi x$). Second-order borderlines are then supposed to come between the definitely-definitely- Φ items and the definitely first-order borderlines, being neither one nor the other ($\sim\text{Def}\text{Def}\Phi x$ & $\sim\text{Def}\sim\text{Def}\Phi x$). It would seem to follow, then, that if definite first-order borderlines are impossible—if there are only plain old regular first-order borderlines—any second-order borderlines must instead come between the definitely-definitely- Φ items and the plain old regular first-order borderlines. (There is nothing else for them to come between.) In other words, they must be neither-definitely-definitely- Φ -nor-borderline- Φ ($\sim\text{Def}\text{Def}\Phi x$ & $\sim(\sim\text{Def}\Phi x$ & $\sim\text{Def}\sim\Phi x)$). Then, finally, since every item in the series is not-definitely-borderline, the second-order borderlines must be neither-definitely-definitely- Φ -nor-borderline- Φ -nor-definitely-borderline- Φ ($\sim\text{Def}\text{Def}\Phi x$ & $\sim(\sim\text{Def}\Phi x$ & $\sim\text{Def}\sim\Phi x)$ & $\sim\text{Def}(\sim\text{Def}\Phi x$ & $\sim\text{Def}\sim\Phi x)$). Figure 5 provides an illustration, using ‘tall’:

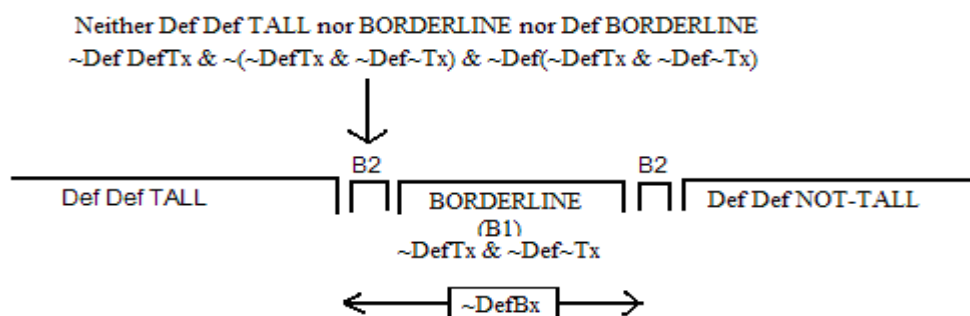


Figure 5

The trouble of course is that anything that is not-borderline is either definitely tall or definitely not-tall, which is incompatible with its being a borderline case of any order. Contrary to the objection, first-order borderlines are not second-order borderlines. Nothing can be second-order borderline.

Perhaps we have thought that vague predicates could have higher-order borderline cases because we allowed technique to lead intuition: we were enchanted by the formal permissibility of generating certain expressions with the definiteness operator. What is possible, even coherent, in natural language may be a different matter. My own view (2005), which I will not elaborate here, is that borderline cases are not properly defined using a definiteness operator.⁸ In order to connect with the philosophical literature I have been going along with the standard definition; but in what follows I will no longer use that device. (What I want to say will not commit us to any particular analysis of borderline cases.)

2. Prescriptive Higher-Order Vagueness

Prescriptive higher-order vagueness appears to be a feature of certain metalinguistic predicates, such as ‘mandates application of “ Φ ”’ and ‘can competently be called “ Φ ”’, that have to do with

the proper application of a vague word. I think the vagueness of these predicates has been misunderstood. (Actually, their being metalinguistic is probably inessential to the view I'll sketch below; my argument may apply equally to 'mandates being classified as Φ ' and 'can competently be judged Φ ', for example. I do not know how to distinguish the relevant family of terms in a principled way, but that should not cause trouble for us here.)

To begin, consider the vague predicate 'old'. Presumably a hundred-year-old person mandates application of 'old' (for a person): failing to call him 'old' would be linguistically unacceptable, incompetent, a mistake. In contrast, a 63-year-old person can be classified as old, but can also be classified as borderline old or as (e.g.) middle-aged. Different competent speakers, and each competent speaker on different occasions, may classify the 63-year-old differently. Similarly, a 50-year-old person mandates classification as old for a ballet dancer, whereas a 35-year-old could be called 'old' or 'borderline old' or 'middle-aged' for a dancer. I will say that cases like the 63-year-old and the 35-year-old permit variable classification relative to the specified comparison class:⁹ we are free to apply the predicate 'old' and also free to withhold it. By the same token, variability of application in the neighborhood of its blurred boundaries is characteristic of—indeed, I would argue, essential to—competent use of a vague predicate. In a sorites series this variability is reflected in a multiplicity of permissible stopping places. In a sorites series of ages proceeding from 100 years to 1 year by increments of one year, on a given occasion you might stop applying 'old' (for a person) at 70 whereas I stop at 65. And you might stop at 69 the next time around. There is no question of error, because our particular stopping places are arbitrary: in every case we could as easily, as competently, have stopped elsewhere.¹⁰ To put the point another way, there is no reason, in the nature of the case, to shift at any particular place.

If the predicate ‘mandates “old”’ is vague, it too should permit arbitrarily variable application. But see what happens when ‘mandates “old”’ is applied to the series of ages. Suppose that on a given occasion you stop applying ‘mandates “old”’ at (after) 70. You stop there arbitrarily; indeed let us suppose that, as must typically be the case, you are well aware that your stopping place is arbitrary. To suppose that your stopping place is arbitrary is to suppose that ‘mandates “old”’ could also permissibly be withheld from 70. But as a moment’s reflection reveals, if it is permissible to withhold ‘mandates “old”’ from 70, it is permissible to withhold ‘old’ from 70. Hence you find yourself in the incoherent position of saying that 70 mandates application of ‘old’ while also granting that ‘old’ can permissibly be withheld from 70. It is one thing to judge that 70 is old while granting that it’s permissible to withhold ‘old’ from 70, and quite another thing to judge that 70 *mandates* ‘old’ while granting that it’s permissible to withhold ‘old’ from 70. The relevant difference is that the latter case involves a legislative judgment, *viz.*, the judgment that 70 must be called ‘old’ on pain of incompetence. And the trouble is that the legislative force of ‘mandates’ and the permissible variability of the application of ‘old’ pull in opposite directions. The predicate ‘mandates application of “old”’, which brings the two together, appears internally conflicted.

An analogous difficulty seems to arise for the prescriptive metalinguistic predicate ‘can competently be called “old”’. Suppose I stop applying ‘can competently be called “old”’ (for a person) at (after) 50 years. I know that I stop there arbitrarily, which is to say that I grant the permissibility of withholding the predicate ‘can competently be called “old”’ from 50. But now consider the speaker who chooses to withhold the latter predicate. She must grant the permissibility of applying it. But as a moment’s reflection reveals, if it is permissible to apply ‘can competently be called “old”’ to 50, then 50 can competently be called ‘old’. Thus the speaker ends

up in the incoherent position of supposing that 50 can competently be called ‘old’ while herself withholding the predicate ‘can competently be called “old”’ from 50. It is one thing to withhold ‘old’ from 50 while granting that 50 can competently be called ‘old’, and quite another thing to withhold ‘can *competently* be called “old”’ from 50 while granting that 50 can competently be called ‘old’. The legislative character of ‘can competently be called “old”’ is what underwrites this distinction.

The preceding discussion suggests that there can be no arbitrary permissible stopping places in a sorites series for ‘mandates “old”’ or ‘can competently be called “old”’. (This seems to me independently plausible. How could what is mandatory or competent in the English language vary, arbitrarily, from speaker to speaker and time to time?) At the same time, there can be no permissible *non*-arbitrary stopping places: in other words, there can be no sharp boundary between the ages that mandate application of ‘old’ and the ages that don’t, or between the ages that can competently be called ‘old’ and the ages that cannot. It seems to follow, then, that there can be no permissible stopping places at all in sorites series for these legislative predicates. This is a baffling result, to say the least. How should we respond to it?

The solution, I think, is to recognize that although the surface grammar of ‘mandates “old”’ (*simile* ‘can competently be called “old”’) has it applying to chronological ages, the predicate also, implicitly, makes reference to the verbal behavior of users of its embedded vague term ‘old’. In order to judge whether a given age mandates ‘old’, you need to know not only its number of years, but also how other speakers would classify it. A crude initial proposal might be this: a given age n mandates application of ‘old’ just in case, on average, almost all competent English speakers would apply ‘old’ to n .¹¹ (By ‘competent English speakers’ I mean only that they are generally competent at speaking English; the question whether they are competent spe-

cifically in the use of ‘old’ remains open for the moment.) What this means—here is the point that has been missed—is that a sorites series for ‘old’ is not a sorites series for ‘mandates “old”’. Instead, a sorites series for ‘mandates “old”’ may be a series of pairs, each containing a chronological age together with an average percentage of competent English speakers who would apply ‘old’ to that age.¹² Figure 5 illustrates such a series beginning with an age of 100 years together with an average of 99 percent of competent speakers applying ‘old’, progressing to an age of 1 year together with an average of 0.5 percent of speakers applying ‘old’:




Fig.5. Sorites series for ‘mandates application of “old”’.

On a given occasion, if you proceeded along this series, you might stop applying ‘mandates “old”’ at 97 percent, while I might stop at 95 percent. And you might stop at 90 percent the next time around. (Stopping at 97 percent would commit you to supposing that, on average, 3 percent of competent English speakers would use the word ‘old’ incompetently at any given time.)

When ‘mandates “old”’ is applied to the right kind of sorites series, we can see that, *qua* vague predicate, it behaves in the same manner as the lexical predicate ‘old’. Its competent application is arbitrarily variable. If you stop applying ‘mandates “old”’ at 97 percent, you do so, and know that you do so, arbitrarily: you could as easily have stopped at 97.5 percent for example. No incoherence results, because you are no longer making a legislative judgment; you are making a merely descriptive judgment as to whether application of ‘old’ by 97 percent of competent English speakers is sufficient to make it the case that the corresponding age—whatever it may be—mandates application of ‘old’. As far as I can see, there is nothing incoherent about

judging that 97 percent makes a given age mandate ‘old’, while also granting the permissibility of withholding that judgment. You and I may permissibly vary in our judgments as to whether a given percentage makes application of a predicate mandatory.

Perhaps it seems that we can say “categorically”, without considering anyone’s verbal behavior, that 100 years mandates application of ‘old’. (Maybe the ages down through about 70 seem this way.) Intuitively: 100 years is old no matter what anyone else says. Similarly, it may seem that we can say, without knowing anything about anyone’s verbal behavior, that a pure blue patch mandates application of ‘blue’. It may seem that the predicate ‘mandates “blue”’ can competently be applied to such a patch *upon inspection alone, just by looking*. Judgments in very central cases do not seem to rest upon consideration of anyone’s verbal behavior. However, where an item is not a highly central case, a judgment as to whether it mandates application of a certain predicate, or whether a certain predicate can competently be applied to it, may be impossible apart from some knowledge of what other competent speakers would say. We can apply ‘blue’ just by looking, but not ‘mandates “blue”’; we can apply ‘old’ just by considering the number of years, but not ‘mandates “old”’.

Our discussion of the prescriptive predicates suggests that in at least some cases, we cannot add conditions to the application of a vague predicate (e.g., that it be mandatory) without thereby generating a new vague predicate (e.g., ‘mandates “old”’) requiring a new sorites series. To see what can happen if we overlook this rule, consider the following passage from Timothy Williamson’s influential book *Vagueness*:

On the view that nothing is hidden [in particular a semantic view—DR], it should be harmless to imagine omniscient speakers, ignorant of nothing relevant to the borderline case... Accompanied by an omniscient speaker of Eng-

lish, you remove grain after grain from a heap. After each removal you ask ‘Is there still a heap?’. ... For some number n , she says ‘Yes’ after each of the first n removals, but not after $n + 1$... You repeat the experiment with other omniscient speakers.... If they all stop at the same point, it evidently does mark some sort of previously hidden boundary.... [A non-epistemic view] must therefore hold that different omniscient speakers would stop at different points. They are conceived as having some sort of discretion...

You can instruct the omniscient speakers... to use their discretion... conservatively, so that they answer ‘Yes’ to as few questions as is permissible. ... Now if two omniscient speakers stop answering ‘Yes’ at different points, both having been instructed to be conservative, the one who stops later has disobeyed your instructions, for the actions of the other show that the former could have used her discretion to answer ‘Yes’ to fewer questions than she actually did. But the omniscient speakers are cooperative. They will... obey your instructions... It is not as though, however many times they said ‘yes’, they could have said it fewer times, for the sorites series is finite... Thus if all [omniscient speakers] are instructed to be conservative, all will stop at the same point. You do not know in advance where it will come.

It marks some sort of previously hidden boundary...(pp.199-200).

Before I say anything about this argument in connection with prescriptive predicates, I want to make sure that its most obvious mistake *is* obvious. The instruction to “answer ‘Yes’ to as few questions as is permissible” is equivalent to an instruction to stop applying ‘heap’ at the earliest (most conservative) permissible place. But the semanticist about vagueness denies that there is

such a place; the instruction cannot be carried out. (According to the semanticist, an omniscient speaker just is a competent speaker.) Only those who already believe in the existence of a sharp cut-off will imagine that it can.

Now to the question of requiring a new sorites series. (We will need to work around the mistake just mentioned.) In instructing the omniscient speakers to stop applying ‘heap’ at the most conservative permissible place, Williamson is in fact asking them to apply a new predicate—something like ‘is a permissible stopping place for the predicate “heap”’.¹³ Our reflections on ‘mandates “old”’ suggest that whether a given collection of grains satisfies the latter predicate does not depend, or does not depend solely, on its number. Rather, the verdict depends also upon the way in which the community of English speakers applies the word ‘heap’. We can imagine a sorites series analogous to the one for ‘mandates “old”’:




Fig. 6. Sorites series for ‘is a permissible stopping place for “heap”’.

When ‘is a permissible stopping place for “heap”’ is applied to the right sort of sorites series, the semanticist can say again that the omniscient speakers (who are, of course, just competent speakers!) have “discretion”, and so may diverge, in their applications of it. A speaker who starts applying ‘is a permissible stopping place for “heap”’ at 90 percent knows that his judgment is arbitrary, so he acknowledges the permissibility of withholding that judgment. His judgment is not legislative, so no incoherence threatens.

3. *Another hierarchy?*

What about the iterative predicates ‘mandates application of “mandates application of ‘old’”’ and ‘mandates application of “mandates application of ‘mandates application of “old’”’’, etc.?

Even supposing we’ve got the right kind of sorites series for each of these expressions, won’t we be stuck with an unending hierarchy of higher-order vague predicates? I don’t think so. To see why, consider again our sorites series for the predicate ‘mandates application of “old’”’ (see the bottom pair of series in Figure 7 below). If you were to proceed along series O, *viz.*, the series of chronological ages, you would be judging whether a given age makes a person old. Proceeding along series MO, you would be judging whether application of ‘old’ by a certain average percentage of ordinary speakers makes it the case that application of ‘old’ to the corresponding age is mandatory.

Consider now a further series, MMO, for the predicate ‘mandates application of “mandates application of ‘old’”’ (‘mandates “mandates ‘old’”’). MMO specifies average percentages of competent English speakers who would apply ‘mandates “old’”’, given percentages of speakers who would apply ‘old’ (as specified in MO). Speakers’ applications of ‘mandates “old’”’ reflect their judgments as to which of the various percentages specified in MO make it the case that application of ‘old’ to the corresponding age is mandatory. More strictly, MMO is a series of triples each containing a percentage of speakers who would apply ‘mandates “old’”’, together with the corresponding percentage from series MO and age from series O. Proceeding along MMO, you would be judging whether application of ‘mandates “old’”’ by a given percentage of speakers makes application of ‘mandates “old’”’ mandatory. As always, you would vary in your judg-

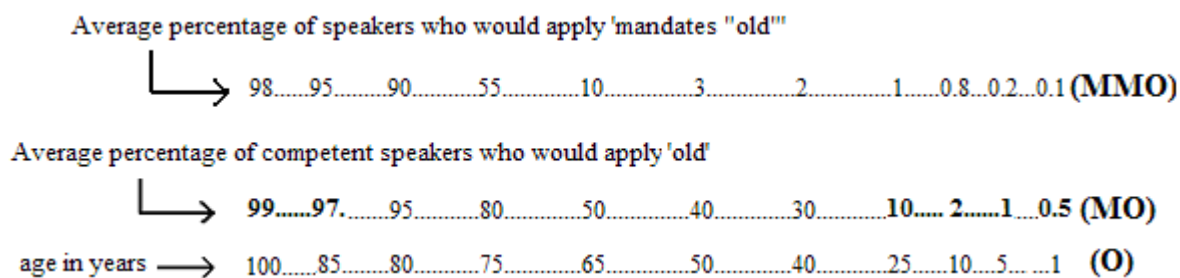


Fig.7. Sorites series (MMO) for 'mandates application of "mandates application of "old"'

ments from occasion to occasion. You might stop applying 'mandates "mandates 'old"' at 95 percent one time, and at 90 percent the next. Analogously for further iterations ("higher orders").

Now the scheme pictured in Figure 7 may appear hierarchical.¹⁴ But consider that, starting with the initial "metalinguistic" series MO for 'mandates "old"', you would always be judging whether application by a given percentage of English speakers makes it the case that application of the predicate in question is mandatory. As suggested in Figure 7, your judgments would vary from series to series (predicate to predicate), but in every series your classifications would depend upon answering the same question: does application of 'Φ' by a certain percentage of speakers show that that application is mandatory? One possibility, then, is that the variations in your classifications of the items (average percentages) in these different series (O, MO, MMO) are just the variations that would occur were you to make repeated runs along any one of them; in particular, your variations across these series may be just the variations that would occur over repeated runs along the MO series. (Perhaps this shows that all of the iterated 'mandates' operators are in effect semantically redundant upon the first.) For this reason I would suggest that, while we can regard the iterative predicates as higher-order vague insofar as they are

metalinguistic and vague, the resulting structure may be better conceived as *recurrent* rather than hierarchical.

In this connection it is interesting to note that whereas ‘mandates “old”’ seems vague, ‘mandates “prime number”’ and ‘mandates “richer than \$110,000”’ seem precise. ‘Mandates “blue”’ seems vague while ‘mandates “6ft tall”’ seems precise. In general, ‘mandates “ Φ ”’ seems vague just in case ‘ Φ ’ is vague, and precise just in case ‘ Φ ’ is precise. The same appears true for the predicate ‘can competently be called “ Φ ”’: it seems vague just in case ‘ Φ ’ is vague. In view of this duality I suggest that, in these metalinguistic uses anyway, the terms ‘mandates’ and ‘competently’ are neither vague nor precise; if ‘mandates “ Φ ”’ and ‘can competently be called “ Φ ”’ are vague, their vagueness must derive entirely from ‘ Φ ’. (I think we can say that the vagueness of ‘mandates “ Φ ”’ and ‘can competently be called “ Φ ”’ just is the vagueness of ‘ Φ ’.) The distinction between vagueness and precision is often taken to be exhaustive for predicates, but I know of no good reason why. In fact, Russell says that ‘vague’ and ‘precise’ are contraries, not contradictories: “We are able to conceive precision; indeed, if we could not do so, we could not conceive vagueness, which is merely the contrary of precision” (1999:65).

The point I want to make is that if ‘mandates’ is not itself vague, then iterating it, as in ‘mandates ‘mandates “ Φ ”’ and ‘mandates “mandates ‘mandates “ Φ ””’, etc., does not introduce any additional vagueness. Intuitively, the predicate ‘mandates “mandates ‘mandates “ Φ ””’ is no vaguer than ‘mandates “ Φ ”’, which in turn is no vaguer than ‘ Φ ’. This result lends further credence to the idea, floated above, that iterated ‘mandates’ operators are semantically redundant. And if that is right, then even if the structure pictured in Figure 7 is a hierarchy of some unobvious sort, it is not a hierarchy of vaguenesses.

Of course, 'mandates' and 'competent' and their ilk have other uses, including object-linguistic ones, that may differ significantly from the metalinguistic uses we have been discussing. For instance, actions can be mandatory or permissible, and doctors and teachers can be competent. Of particular relevance to an understanding of vagueness is the fact that speakers may or may not be competent. Whether any conclusions drawn here about the metalinguistic uses of these words will transfer to their object-linguistic uses is a matter I leave for further investigation.

Notes

¹ See e.g. Wright, this volume. A third form of higher-order vagueness is the generic (i.e., not necessarily prescriptive) vagueness of the metalanguage in which a theory of vagueness is formulated, where that metalanguage is a natural language like English. For example, perhaps vagueness is defined in terms of a certain kind of context-relativity. The word ‘context’ is probably vague. I take for granted that this kind of higher-order vagueness exists.

² For some of the important work that has been done on higher-order vagueness, see Deas 1989, Graff 2003, Heck 1993, Hyde 1994 and 2003, Sorensen 1995, Tye 1994, Varzi 2003, and Wright 1992, 1994, and this volume.

³ The hyphenation is only to avoid scope ambiguities.

⁴ See my 2005, note 18; also Graff 2003.

⁵ I will underline when it is convenient to refer to the category (type, kind, property, class) named by a vague predicate, rather than to the predicate itself.

⁶ Here of course I refer to a hypothetical competent, sincere, cooperative speaker who fails to apply ‘ Φ ’ upon being queried. Feel free to add whatever further specifications you think necessary for present purposes.

⁷ An anonymous reviewer writes: “if I am [instructed] to classify various colours on the red-yellow continuum as either red, yellow or borderline then, if I failed to classify an orange patch as borderline, wouldn't I be mistaken?” I don't see why. If the patch *is orange*, how could you be mistaken in failing to call it ‘borderline’? Unless ‘borderline’ just *means* ‘orange’ (in which case we are not speaking English), the correct response to the instruction is to say that it cannot

be carried out because it fails to provide adequate response categories. By the same token, if you were instructed to classify colors on the “red-green continuum” as either red, green, or borderline, you could hardly be convicted of error if you failed to classify a yellow patch as borderline. Such an instruction would be illegitimate (incompetent, if you like). Rejecting an instruction is not the same as making a mistake.

Here’s another way to put the point. If you were required to call the orange patch ‘borderline’, then you would be using ‘borderline’ to name what would in fact be a non-borderline category—a category in its own right on a par with red and yellow, namely orange. In other words, you would be using ‘borderline’ to mean ‘orange’. See the discussion on p.4 above.

A second anonymous reviewer writes that (ii) in the Simple Argument needs to be restricted to avoid counterexamples based on an abuse of privilege. If I persistently characterize what you regard as borderline cases as clear cases, then I have at least abused a right. Think of prospective employees puffing their credentials.

I don’t see why such cases would have to be classified as borderline. Couldn’t they fairly be classified as (definitely) not Φ ?

⁸ I propose an analysis of ‘borderline case’ that is bivalent, does not employ a definiteness operator, and eliminates the possibility of higher-order borderlines. On my view, borderlines are properly defined in terms of contrary or incompatible predicates, such as ‘old’ and ‘young’, or ‘old’ and ‘middle-aged’, rather than contradictory ones like ‘old’ and ‘not-old’.

⁹ Make the comparison class or context as fine-grained as you like; the variability will persist. I elaborate, and provide some experimental evidence, in my 2008.

¹⁰ Consider that if our stopping places are not arbitrary, if reasons or argument can be given for stopping at one place rather than another, then the increments in the series are not small enough for a *sorites* series.

¹¹ This crude proposal is doubtless incorrect, but it has the virtue of simplicity. Perhaps, for example, it would be better to say that a given age n mandates application of ‘old’ just in case, on average, almost all competent English speakers would respond to a failure to apply ‘old’ to n by, say, expressing bafflement, or correcting the speaker in question, or initiating an argument with the speaker in question.

¹² Presumably, as a matter of empirical fact, the percentage of competent English speakers who would apply ‘old’ to a given age (when queried, etc.) varies across time. Hence the reference to an average percentage.

¹³ Actually, the predicate tacitly introduced by Williamson is ‘*earliest* permissible stopping place...’; but as I said, we must work around this mistake.

¹⁴ It may also appear mind-numbingly complex; but that owes to the mind-numbing complexity of the predicates at issue. No one ever actually uses such crazy words.

Bibliography

- Deas, R., 1989. "Sorensen's Sorites". *Analysis* 49: 26-31.
- Graff, D., 2003. "Gap Principles, Penumbral Consequence, and Infinitely Higher-Order Vagueness". In *Liars and Heaps: New Essays on Paradox*, ed. J. C. Beall (Oxford: Oxford University Press), pp.195-222.
- Heck, R., 1993. "A Note on the Logic of (Higher-Order) Vagueness". *Analysis* 53, pp.201-8.
- Hyde, D., 1994. "Why Higher-Order Vagueness is a Pseudo-Problem". *Mind* 103: 35-41.
- Hyde, D., 2003. "Higher Orders of Vagueness Reinstated". *Mind* 112, 46: 301-5.
- Keefe, R., and Smith, P. (1999). *Vagueness: A Reader* (MIT)
- Lindsey, D., Raffman, D., and Brown, A., 2008 (in progress). "Psychological Hysteresis and the Nontransitivity of Insignificant Differences".
- Raffman, D., 2005. "Borderline Cases and Bivalence". *Philosophical Review* 114, 1.
- Raffman, D., 2008. _____.
- Russell, B., 1923. "Vagueness". *Australasian Journal of Philosophy and Psychology* 1, 84-92.
Reprinted in Keefe and Smith 1999, 61-68.
- Sainsbury, R.M., 1997 (1990). "Concepts Without Boundaries". Orig. Inaugural Lecture, King's College London, 1990. Reprinted in Rosanna Keefe & Peter Smith (eds.), *A Vagueness Reader* (MIT,1997), pp.251-264.
- Shapiro, S., 2006. *Vagueness in Context*. Oxford.
- Sorensen, R., 1985. "An Argument for the Vagueness of 'Vague'". *Analysis* 45: 134-37.
- Tye, M., 1994. "Why the Vague Need *Not* Be Higher-Order Vague". *Mind* 103: 43-5.

Varzi, A., 2003. "Higher-Order Vagueness and the Vagueness of 'Vague'". *Mind* 112: 295-99.

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

Wright, C., 1992. "Is Higher-Order Vagueness Coherent?". *Analysis* 52: 129-39.

Wright, C., 1994. "The Epistemic Conception of Vagueness". *Southern Journal of Philosophy*

33, Spindel Supplement, pp.133-59.