

How vagueness could cut out at any order

Timothy Williamson (1999) presents (without endorsing) an argument that for any $n > 1$, any sentence that is n th order vague is $(n + 1)$ th order vague. The setting for the argument is a standard, Kripke-style model theory for a propositional language enriched with an operator Δ for ‘Definitely’. A *frame* is an ordered pair $\langle W, < \rangle$, where W is a set (of “points”) and $<$ is a binary relation (of “accessibility”) on that set. A *model* is a triple $\langle W, <, \llbracket \cdot \rrbracket \rangle$, where $\langle W, < \rangle$ is a frame and $\llbracket \cdot \rrbracket$ a function from the sentences of the language to subsets of W , subject to the standard recursive rules: $\llbracket \neg \alpha \rrbracket = W \setminus \llbracket \alpha \rrbracket$; $\llbracket \alpha \vee \gamma \rrbracket = \llbracket \alpha \rrbracket \cup \llbracket \gamma \rrbracket$; $\llbracket \Delta \alpha \rrbracket = \{x : \forall y (x < y \rightarrow y \in \llbracket \alpha \rrbracket)\}$. A sentence α is *precise* at $x \in W$ iff $x \in \llbracket \Delta \alpha \vee \Delta \neg \alpha \rrbracket$; otherwise, α is *vague* at x . α is *precise in a model* iff α is precise at every point in the model, otherwise α is *vague in the model*. (Note that in a *connected* model, i.e. one where any proper subset of W contains some point from which some point not belonging to that subset is accessible, α is precise in the model iff $\llbracket \alpha \rrbracket = W$ or $\llbracket \alpha \rrbracket = \emptyset$.)¹

Williamson proposes the following as the most natural definition of higher-order vagueness. For any set X of sentences, let $\mathbf{C}X$ be the closure of X under the truth-functional operators (\neg and \vee), and $\mathbf{C}_n X$ is defined inductively as follows:

$$\mathbf{C}_1 X = \mathbf{C}X$$

$$\mathbf{C}_{n+1} X = \mathbf{C} \mathbf{D} \mathbf{C}_n X$$

α is n th-order precise at x iff every member of $\mathbf{C}_n \{\alpha\}$ is precise at x ; otherwise, α is

¹How do these notions relate to our ordinary, non-model-relative notions of vagueness and precision? We clearly want to allow that a vague sentence (“Obama is bald”) might happen to be definitely true or definitely false at the actual world. Williamson suggests that a sentence α is precise iff it is “semantically guaranteed” that $\Delta \alpha \vee \neg \Delta \alpha$ is true. He also suggests that this notion of “semantic guarantee” can be understood model-theoretically: there is some privileged class of models such that a sentence α is semantically guaranteed to be true iff $\llbracket \alpha \rrbracket = W$ in each of them. This is a demanding status: when the truth of α is semantically guaranteed, the truth of $\Delta^n \alpha$ is also semantically guaranteed, for every n . The argument in [self-citation deleted for blind review] suggests that it is too demanding to be interesting. But for present purposes, it doesn’t matter what relation we might take there to be between the defined notions of vagueness-at-a-point and vagueness-in-a-model and vagueness itself.

n th-order vague at x . α is n th-order precise in a model iff α is n th-order precise at every point in its domain; otherwise, α is n th-order vague in the model. Notice that whereas n th-order precision at a point doesn't entail $(n + 1)$ th-order precision at the point, n th-order precision in a model does entail $(n + 1)$ th-order precision in the model.

Say that n is the *cutoff* of α in a model iff α is n th-order vague but $n + 1$ th-order precise in that model. And say that n is the *maximum cutoff* of a class of models iff some sentence has cutoff n in some model in the class, but no sentence has cutoff $> n$ in any model in the class. Williamson (1999) shows that any class of models validating the schema

$$\text{B} \quad \neg\Delta\neg\Delta\alpha \rightarrow \alpha,$$

has maximum cutoff 1. That is: in each such model, every sentence is either precise, first-order vague but second-order precise, or vague at all orders. The proof is simple. Suppose α is n th-order vague for some $n > 1$. Then some member of $\mathbb{C}_n\alpha$ is vague. If so, $\Delta\beta$ must be vague for some $\beta \in \mathbb{C}_{n-1}\alpha$, since $\mathbb{C}_n\{\alpha\}$ is just the closure of these sentences under the precision-preserving operations of negation and conjunction. But given B, $\Delta\beta$ is logically equivalent to $\Delta\neg\Delta\neg\Delta\beta$, so the latter, which belongs to $\mathbb{C}_{n+2}\{\alpha\}$, is also vague; so α is $n + 2$ th-order vague, and hence also $n + 1$ th-order vague.

If B were in fact valid, this result would be illuminating: it would suggest that the classification of sentences according to higher-order vagueness has a kind of underlying simplicity that one might not have suspected. But is B valid? One should not expect to be able to answer this question without considering foundational issues concerning the nature of vagueness and the meaning of Δ . In the present paper, our main concern will be with the epistemicist account of the nature of vagueness developed in Williamson 1994. Should a Williamsonian epistemicist accept B? Williamson (1999) suggests, tentatively, that the answer is yes, on the grounds that his version of epistemicism can be understood ('to a first approximation') as the view that "definiteness is truth under all interpretations of the language indiscriminable from the correct one", and that indiscriminability is a symmetric relation (Williamson 1990, pp. 10–11).

Williamson expresses some misgivings about the adequacy of this gloss on his

view, but let us set these aside.² Still, the argument from the view to the validity of B is not straightforward. To evaluate it, we will need to consider what upshot the proposed account of *definiteness* (as a property of sentences) has for the truth-conditions of sentences involving the ‘definitely’ operator. Certainly $\Delta\alpha$ is supposed to be true (i.e. true on the correct interpretation) if and only if α is definite. So the proposal entails (E):

(E) For any sentence α , $\Delta\alpha$ is true on the correct interpretation iff α is true on all interpretations that are indiscriminable from the correct interpretation.

This does not yet tell us all that much about the meaning of Δ , since it does not tell us about the truth value of $\Delta\alpha$ at worlds other than the actual world. But there are two natural ways of extending (E) to provide such truth-conditions:

(E₁) For any sentence α and world w , $\Delta\alpha$ is true at w on the correct interpretation iff α is true at w on all interpretations that are indiscriminable from the correct interpretation.

(E₂) For any sentence α and world w , $\Delta\alpha$ is true at w on the correct interpretation iff α is true at w on all interpretations that are indiscriminable at w from the correct interpretation.

Here ‘the correct interpretation’ always means the interpretation that is *in fact* correct, not the interpretation that is correct at w ; the difference has to do whether the relation of indiscriminability has to hold at the actual world or at the relevant world w . This would make no difference if facts about which pairs of interpretations are indiscriminable were all necessary; but presumably they are not, since our powers of discrimination are clearly contingent. It is not clear which of (E₁) and (E₂) epistemicists should prefer, insofar as they accept (E), it is plausible that they should at least be prepared to assert their disjunction.

²Caie (2012) raises one major concern: he argues that “‘Everest’ refers to Everest’ is not true under all interpretations indiscriminable from the correct one—since some such interpretations assign a different referent to ‘Everest’ while still assigning the same intension to ‘refers’—and hence will be wrongly counted as indefinite on the proposal.

It might seem that we could derive B from (E) and the symmetry of indiscriminability, as follows. Suppose that $\neg\Delta\neg\Delta\alpha$ is true on the correct interpretation, which we will dub 'c'. Then $\Delta\neg\Delta\alpha$ is not true on c. So by (E), it is not the case that $\neg\Delta\alpha$ is true on all interpretations indiscriminable from c. Since all interpretations indiscriminable from c are bivalent, it follows that there is an interpretation *i* indiscriminable from c on which $\Delta\alpha$ is true. So by (E) again, α is true on every interpretation indiscriminable from *i*. But by the symmetry of indiscriminability, c is indiscriminable from *i*. Therefore α is true on c.

There is a big hole in this argument at the point where we appeal to (E) for the second time, to go from the truth of $\Delta\alpha$ on *i* to the truth of α on all interpretations indiscriminable from *i*. For (E) merely tells us about the conditions under which $\Delta\alpha$ is *true*, i.e. true on the correct interpretation. It says nothing about the conditions under which $\Delta\alpha$ is true on interpretations other than the correct one.

We could plug this hole by strengthening (E) to (E+):

(E+) For any interpretation *i* that is indiscriminable from the correct one and any sentence α , $\Delta\alpha$ is true on *i* iff α is true on all interpretations indiscriminable from *i*.

This suffices to run the argument, since the relevant interpretation *i* is by hypothesis indiscriminable from c.³

But does Williamsonian epistemicism support (E+)? Mahtani (2008) argues that it does not, and that it in fact supports the claim that there are counterexamples to B. One way of developing Mahtani's central point is to appeal to the evident fact that 'indiscriminable' is vague.⁴ Thus, if we are willing to believe in interpretations

³Note that it is not on the cards to say that $\Delta\alpha$ is true on an *arbitrary* interpretation iff α is true on all interpretations indiscriminable from that interpretation. An interpretation is just a function from linguistic items (considered as not having their semantic properties essentially) to semantic values. Whatever semantic role ' Δ ' plays on the correct interpretation, there will of course be incorrect interpretations which treat ' Δ ' in some very different way—e.g. by treating it in the way the correct interpretation treats ' \neg '.

⁴Mahtani, instead, appeals to vagueness in the word 'definitely'. I think we are on somewhat safer ground in appealing to the vagueness of 'indiscriminable', since it is

indiscriminable from c on which ‘bald’ expresses a property that some bald things lack, we should be equally willing to believe in interpretations indiscriminable from c on which ‘indiscriminable’ expresses a relation that fails to hold between some indiscriminable things. Moreover, if such variant interpretations exist, one would need a very special argument to rule out the existence of an interpretation i with the following features:

- (i) i is correct at some world v close to the actual world, @.
- (ii) i is indiscriminable from the correct interpretation c , both at v and at @.
- (iii) i maps ‘indiscriminable’ to a relation R that i does not bear to c , either at w or at @.
- (iv) i is exactly like c as regards the interpretation of all vocabulary other than ‘indiscriminable’, ‘ Δ ’, ‘definite’, etc.
- (v) Some sentence α is false at @ and w according to c , but true at @ and w on all interpretations that bear R to i at either @ or at v .

(i), (iv), and (v) together entail that the open sentences ‘ α is true at w on all interpretations that are indiscriminable from the correct interpretation’ and ‘ α is true at w on all interpretations that are indiscriminable at w from the correct interpretation’ are both true on i at v , relative to an assignment that maps the variable w to @. So if either (E₁) or (E₂) were true on i at v , it would follow that the open sentence ‘ $\Delta\alpha$ is true at w on the correct interpretation’ is also true on i at v relative to such an assignment. If so, given (iv) and the fact that i is the correct interpretation at v , $\Delta\alpha$ must be true on i at @, and so $\neg\Delta\alpha$ must be false on i at @. But in that case, given (E) and the fact that i is indiscriminable from the correct interpretation, $\neg\Delta\neg\Delta\alpha$ must be true at @ on c , the interpretation that is in fact correct. Since α is false on c , we have a counterexample to B.

less clear what kind of semantic value we should assign to ‘definitely’.

So, preserving B requires either denying that there is an interpretation meeting the above conditions—which looks like a tall order if ‘indiscriminable’ is vague in the usual way—or being open to the possibility that there is a close world such that neither (E₁) nor (E₂) is true on the interpretation that is correct at that world. But given Williamson’s thesis that “metalinguistic safety” is a necessary condition on knowledge—a thesis that plays a central role in his overall view—the conclusion that there is such a close world would entail that in the actual world, those who utter (E₁), (E₂), or their disjunction do not thereby express knowledge.⁵ But such a conclusion would be intolerable. By Williamson’s own lights (Williamson 2000, ch. 11), in asserting his theory of vagueness, he is implicitly committing himself to the claim that he is speaking knowledgeably in doing so.⁶

Thus, Williamson’s strategy for arguing that sentences that are vague at any higher order are vague at all higher orders breaks down. But of course, that is not yet a positive reason for the Williamsonian epistemicist to think that there *are* sentences whose vagueness cuts out at finite orders greater than one.

To make progress with the question how the Williamsonian epistemicist should expect the different orders of vagueness to be distributed, it is natural to look for a

⁵For a careful examination of this thesis, see Kearns and Magidor 2008.

⁶In conversation, Williamson has suggested that although ‘ Δ ’ is vague, there are some purposes for which we should want to abstract away from its vagueness. His idea is that the vagueness in $\Delta\alpha$ should be regarded as a product of two separate factors: the vagueness of Δ , and something distinctive to α . When we want to isolate the latter factor, we can do so by using an artificially circumscribed version of Δ , which only takes into account interpretations which agree with the correct interpretation as regards the interpretation of Δ . The obvious way to introduce this new version of Δ is to stipulate that $\Delta\alpha$ is true iff α is true on each interpretation that is indiscriminable from the correct interpretation and also interprets ‘indiscriminable’ in exactly the same way as the correct interpretation. But it is not clear that there is any insight to be gained by the use of an operator that singles out the word ‘indiscriminable’ for special treatment in this way—there don’t seem to be philosophically interesting differences between the way ‘indiscriminable’ is vague and the way any other word is vague. And it is not clear that there is any other way to make sense of the proposal to evaluate sentences involving iterated Δ s in a way that abstracts away from the vagueness of Δ itself. Prima facie, the idea of such abstraction is no clearer than the proposal to introduce an artificially circumscribed use of ‘know’ which somehow abstracts away from the role that one’s ignorance about knowledge in general, as opposed to some factor specific to α , can play in making ‘I don’t know whether I know that α ’ true.

model-theoretic implementation of the way in which differences in the interpretation of words like ‘indiscriminable’ make it easier for $\Delta\alpha$ to be true on some interpretations than on others. Mahtani suggests a natural class of frames which capture this idea in a simplified and tractable way. In what I will call a ‘variable radius frame’, the accessibility relation is determined by a metric d on W , together with a function r which assigns each point to a positive real number, its ‘accessibility radius’:

$$x < y \text{ iff } d(x, y) \leq r(x).^7$$

Counterexamples to the symmetry of $<$ in a variable radius frame will thus occur whenever $r(y) \leq d(x, y) < r(x)$.⁸

Surprisingly, having introduced the class of variable radius frames, Mahtani does not actually investigate whether there models based on such frames with sentences whose vagueness cuts out at orders higher than 1. But we can answer this question by noting the following fact: *Any countable reflexive frame with no one-way cycles is a variable radius frame.* Here a “one-way cycle” is a finite sequence of points x_0, x_1, \dots, x_n such that $x_0 \ll x_1 \ll \dots \ll x_n \ll x_0$, where $x \ll y$ means that $x < y$ and $y \not< x$.

(*Proof:* Let R be the transitive and reflexive closure of \ll . In a frame without one-way cycles, R is a partial order on W . Let R' be some total order on W that extends R . Since W is countable, we can find some function r from W to the open interval $(1,2)$ such that $r(x) \leq r(y)$ exactly when $xR'y$. Define d as follows:

$$d(x, y) = \begin{cases} 0 & \text{if } x = y \\ 1 & \text{if } x \neq y \text{ and } x < y \text{ and } y < x \\ 2 & \text{if } x \not< y \text{ and } y \not< x \\ (r(y) + r(x))/2 & \text{otherwise} \end{cases}$$

⁷Another very similar class of frames which does an equally good job of capturing the underlying idea, we require instead $x < y$ iff $d(x, y) < r(x)$.

⁸Variable radius models should not be confused with the “variable margin models” introduced in the appendix to Williamson 1994. Variable radius models are just standard Kripke models whose accessibility relation is generated in a particular way, whereas variable margin models are not Kripke models at all.




Figure 1

It is easy to see that $d(x, y) \leq r(x)$ iff $x < y$: in the last case, this is because our choice of r ensures that $r(x) > r(y)$ whenever $x \ll y$. It remains to show that d is a metric. Clearly we have $d(x, y) = 0$ iff $x = y$ and $d(x, y) = d(y, x)$ for all x, y . And d also obeys the triangle inequality, since $d(x, y) + d(y, z) \geq 2 \geq d(x, z)$ whenever $x \neq y$ and $y \neq z$, while $d(x, y) + d(y, z) = d(x, z)$ whenever $x = y$ or $y = z$.)

Given this sufficient condition, we can see that the class of variable radius frames includes a class of frames from Williamson 1999, in which Williamson has already established that there is no maximum cutoff. The domain in each of these frames consists of the first n positive integers: $x < y$ iff $|x - y| \leq 1$ or $y = 1$. (See figure 1.) Williamson shows that in a model where $\llbracket p \rrbracket = \{n\}$, $\neg p$ is $n - 1$ th-order vague but n th-order precise. (See appendix (i) of Williamson 1999 for proof.) Since these frames are countable, reflexive, and lack one-way cycles, they are variable radius frames.⁹

Nevertheless, as Williamson points out, the structure of the accessibility relation in these frames is odd and unmotivated. And learning that they are variable radius frames does not do very much to advance the question whether vagueness can in fact cut out at higher orders, since it leaves it open that we might identify some further plausible constraints on the behaviour of d and r which would suffice for a maximum cutoff. For example, the following conditions on d and r seem quite natural:

Continuity: For any two points x and y and $r \in [0, 1]$, there is a point z such that

$$d(x, z) = rd(x, y) \text{ and } d(z, y) = (1 - r)d(x, y).$$

Gradualness: Whenever $d(x, y) \leq r(x)$, $1/2 < r(x)/r(y) < 2$.

Continuity looks reasonable on an epistemicist interpretation of the models, given the

⁹For example, we can let $r(1) = 1$, $r(x) = 3/2$ when $x > 1$, $d(x, y) = 0$ when $x = y$, $d(x, y) = 1$ when $|x - y| = 1$, $d(x, y) = 3/2$ when $|x - y| > 1$ and one of x and y is 1, and $d(x, y) = 2$ otherwise.

continuously varying character of the space of possible interpretations. *Gradualness* also looks reasonable, at least if we allow ourselves to restrict our attention to interpretations that are not too radically dissimilar from the correct one. For points x and y whose accessibility radii differed by a factor of 2 or more would correspond to interpretations whose interpretations of ‘indiscriminable’ are *very* different; any interpretation of ‘indiscriminable’ under which such dissimilar points could count as ‘indiscriminable’ would be radically unlike the correct one.

It is easy to see that whenever r and d satisfy *Continuity* and *Gradualness*, the accessibility relation defined by $x < y =_{df} d(x, y) \leq r(x)$ will have the property that whenever $x < y$, there is at least one z such that $y < z$ and $z < x$ —e.g. any z for which $d(x, z) = d(y, z) = d(x, y)/2$. And this condition clearly validates the following weakening of B:

$$B^- \quad \neg\Delta\neg\Delta\alpha \supset \alpha.$$

So, Mahtani’s argument that Williamsonian epistemicists should reject B does not suggest any reasons for doubting B^- . Since certain *other* philosophically interesting consequences of B have turned out also to be consequences of B^- , it is thus worth investigating the question whether the class of variable radius models satisfying *Continuity* and *Gradualness* has a maximum cutoff.¹⁰

In the remainder of this paper, I will show that the answer to this question is no, by constructing a variable radius model \mathfrak{M} which satisfies *Continuity* and *Gradualness* but has no maximum cutoff: for every n , there is a sentence—an atomic sentence, in fact—that is n th-order vague but $n + 1$ th-order precise. Moreover, this model’s metric and accessibility radius function are intuitively quite “well behaved”, so its existence

¹⁰Some such consequences are discussed in Graff 2002. See also Bacon MS, which shows that certain other interesting consequences of B also follow from the following principle:

$$B^* \quad \Delta(\alpha \rightarrow \Delta\alpha) \rightarrow (\neg\alpha \rightarrow \Delta\neg\alpha)$$

Bacon shows that B^* is valid on a frame just in case whenever $x < y$, there is a finite sequence z_1, \dots, z_n , all of whose members are accessible from x , such that $y < z_1$, $z_1 < z_2, \dots$, and $z_n < x$. Like B^- , B^* is valid in any variable-radius model satisfying *Continuity* and *Gradualness*.

tends to suggest that we are unlikely to find further natural structural constraints, in the same spirit as *Continuity* and *Gradualness*, whose imposition would reintroduce a maximum cutoff.

Let's take our object language to have infinitely many atomic sentences $p_0, p_1, p_2 \dots$: the plan is to define \mathfrak{M} in such a way that p_i has cutoff i in \mathfrak{M} . \mathfrak{M} will also be connected, so that α is n th-order precise in \mathfrak{M} iff $\llbracket \mathbf{C}_n\{\alpha\} \rrbracket$ (i.e. $\{\llbracket \beta \rrbracket : \beta \in \mathbf{C}_n\{\alpha\}\}$) is $\{\emptyset, W\}$. So to ensure that p_i has cutoff i , it will suffice to make sure that the model satisfies the following desiderata:

- (1) $\llbracket p_i \rrbracket$ is \emptyset or W if and only if $i = 0$.
- (2) Either $\llbracket \Delta p_{i+1} \rrbracket = \llbracket p_i \rrbracket$ and $\llbracket \Delta \neg p_{i+1} \rrbracket = \emptyset$, or $\llbracket \Delta p_{i+1} \rrbracket = \emptyset$ and $\llbracket \Delta \neg p_{i+1} \rrbracket = \llbracket \neg p_i \rrbracket$.

For given (2), $\llbracket \mathbf{DC}\{p_{i+1}\} \rrbracket$ is either $\{\emptyset, \llbracket p_i \rrbracket, W\}$ or $\{\emptyset, \llbracket \neg p_i \rrbracket, W\}$. In either case, $\llbracket \mathbf{CDC}\{p_{i+1}\} \rrbracket = \{\emptyset, \llbracket p_i \rrbracket, \llbracket \neg p_i \rrbracket, W\} = \llbracket \mathbf{C}\{p_i\} \rrbracket$. Since $\mathbf{C}_{n+1}\{\alpha\}$ is defined as $\mathbf{CDC}_n\{\alpha\}$, a straightforward proof by induction then shows that $\llbracket \mathbf{C}_n\{p_m\} \rrbracket = \llbracket \mathbf{C}\{p_{m+1-n}\} \rrbracket$ for all $n \leq m + 1$. By (1), $\llbracket \mathbf{C}\{p_{m+1-n}\} \rrbracket = \{\emptyset, W\}$ when $n = m + 1$, so that p_m is $m + 1$ th-order precise, but $\llbracket \mathbf{C}\{p_{m+1-n}\} \rrbracket \neq \{\emptyset, W\}$ when $n \leq m$, so that p_m is not precise at any lower order.

That's the plan; now for the implementation. The domain of \mathfrak{M} will be the open interval $(1, 2)$. The metric is the standard one: $d(x, y) = |x - y|$. The accessibility radius function is given by $r(x) = x/3$: thus $x < y$ iff $|x - y| \leq x/3$. This is illustrated in figure 2: the shaded region between the lines $y = 2x/3$ and $y = 4x/3$ corresponds to the set $\{(x, y) : x < y\}$.

$\llbracket p_i \rrbracket$ will always be some open interval $(1, a_i)$. To see what the sequence a_0, a_1, \dots needs to look like, note the following facts (which can be verified by inspection of figure 2, bearing in mind that $\llbracket \Delta \alpha \rrbracket$ is the set of points from which only members of $\llbracket \alpha \rrbracket$ are accessible):

- (i) If $\llbracket \alpha \rrbracket = (1, x)$ where $x < 4/3$, then $\llbracket \Delta \alpha \rrbracket$ is \emptyset and $\llbracket \Delta \neg \alpha \rrbracket$ is $[3x/2, 2)$.
- (ii) If $\llbracket \alpha \rrbracket = (1, x)$ where $x > 4/3$, then $\llbracket \Delta \alpha \rrbracket$ is $(1, 3x/4)$ and $\llbracket \Delta \neg \alpha \rrbracket$ is \emptyset .
- (iii) If $\llbracket \alpha \rrbracket = (1, 4/3)$, then $\llbracket \Delta \alpha \rrbracket = \llbracket \Delta \neg \alpha \rrbracket = \emptyset$.




Figure 2

Let g be the function on $(1, 2)$ such that $g(x) = 3x/2$ if $x \leq 4/3$ and $g(x) = 3x/4$ if $x > 4/3$. Given (i)–(iii), it is enough to satisfy desideratum (2) if we can set things up so that a_i is always equal to $g(a_{i+1})$. And this is easy to do, since g is a bijection from $(1, 2)$ to $(1, 3/2) \cup (3/2, 2]$, and thus has an inverse: $g^{-1}(x) = 4x/3$ if $x \in (1, 3/2)$ and $2x/3$ if $x \in (3/2, 2]$. (The graph of g^{-1} is the union of the two thick lines in figure 2.) So we can generate the sequence a_0, a_1, \dots as follows:

$$a_0 = 2$$

$$a_{i+1} = g^{-1}(a_i)$$

To verify that this yields an infinite sequence, note that, provided a_i exists, it must be a fraction of the form $2^j/3^i$ for some natural number j . Since no such fraction can equal $3/2$, a_i is in the domain of g^{-1} , so a_{i+1} will exist, and we will have $a_i = g(a_{i+1})$ as desired. Moreover, desideratum (1) is clearly also satisfied: $\llbracket p_0 \rrbracket = (1, 2) = W$, and since 2 is not in the range of g^{-1} , $\llbracket p_i \rrbracket$ is a proper subset of W whenever $i > 0$. So, for every n , p_n is n th-order vague but $n + 1$ th-order precise in \mathfrak{M} .

This behaviour of \mathfrak{M} is quite sensitive to the way we chose its domain. Let \mathfrak{M}' be like \mathfrak{M} except that its domain W' is the closed interval $[1, 2]$ rather than the open interval $(1, 2)$. \mathfrak{M}' has a maximum cutoff of 1! For assume for contradiction that some sentence α

has cutoff $n > 1$ in \mathfrak{M} . Then $\llbracket \mathbf{C}_{n+1}\{\alpha\} \rrbracket$ is $\{\emptyset, W'\}$ but $\llbracket \mathbf{C}_n\{\alpha\} \rrbracket$ contains proper subsets of W' . Consider some $\beta \in \mathbf{C}_{n-1}\{\alpha\}$ such that $\llbracket \Delta\beta \rrbracket$ is neither \emptyset nor W' . (There must be some such β given that $\llbracket \mathbf{C}_n\{\alpha\} \rrbracket \neq \{\emptyset, W'\}$.) Then $\llbracket \neg\Delta\beta \rrbracket \neq W'$; so $\llbracket \Delta\neg\Delta\beta \rrbracket \neq W'$; so $\llbracket \Delta\neg\Delta\beta \rrbracket$ must be \emptyset , since $\Delta\neg\Delta\beta$ belongs to $\llbracket \mathbf{C}_{n+1}\{\alpha\} \rrbracket$. If $\llbracket \neg\Delta\beta \rrbracket$ contained $[1, 4/3]$ or $[4/3, 2]$, $\llbracket \Delta\neg\Delta\beta \rrbracket$ would contain 1 or 2 and thus not be empty. So $\llbracket \Delta\beta \rrbracket$ must contain some $x \in [1, 4/3]$ and some $y \in [4/3, 2]$. But in that case, $\llbracket \beta \rrbracket$ must contain both $[1, 4x/3]$ and $[2y/3, 2]$, and so $\llbracket \beta \rrbracket$ must contain all of W' . But in that case we also have $\llbracket \Delta\beta \rrbracket = W'$, contradicting our assumption.¹¹

Thus, our investigation does not support the conclusion that the possibility of high cutoffs is in any sense a *typical* feature of variable radius models satisfying *Continuity* and *Gradualness*. Given that \mathfrak{M} has several quite “artificial” features, it is natural to wonder whether we might find some further natural formal conditions on variable radius frames whose imposition would reintroduce a maximum cutoff (without validating B).

On further investigation, however, many of the artificial features of \mathfrak{M} turn out not to be crucial. For example, \mathfrak{M} is one-dimensional, whereas the space of possible interpretations is presumably very high-dimensional under any reasonable metric. But it is easy to generate an n -dimensional analogue of \mathfrak{M} : let the domain be the set of points in \mathbb{R}^n whose distance from the origin o (on the standard metric) is less than 1, $r(x) = (d(x, o) + 1)/3$, and $\llbracket p_i \rrbracket = \{x : d(x, o) < a_i - 1\}$ (where the a_i are defined as before). This model will meet conditions (1) and (2) for the same reason as \mathfrak{M} .¹²

Another artificial and arguably unrealistic feature of \mathfrak{M} is the fact that r changes *linearly* with distance. But the linearity of r is actually not important: all we really need is for r to be a differentiable function whose derivative is everywhere strictly between

¹¹By similar reasoning, a model like \mathfrak{M} whose domain contains $[1, 2]$ as a subset has a maximum cutoff of 1. On the other hand, we can make the domain of \mathfrak{M} a little bit *smaller* and still have sentences with very high cutoffs. For any n , we can find some $\varepsilon > 0$ such that a restriction of the domain of \mathfrak{M} to $(1 + \varepsilon, 2 - \varepsilon)$ will still contain sentences with all cutoffs up to n : we simply choose ε so as to make sure that $a_i \in (1 + \varepsilon, 2 - \varepsilon)$ whenever $0 < i \leq n$.

¹²Alternatively, we could let r *decrease* with distance from the origin, by letting $r(x) = (2 - d(x, o))/3$ and $\llbracket p_i \rrbracket = \{x : d(x, o) > 2 - a_i\}$.

0 and 1, with $r(3/2) = 1/2$ and $\lim_{x \rightarrow 1} r(x) + \lim_{x \rightarrow 2} r(x) = 1$. Given any such r , we can find an appropriate \mathbb{I} by letting $a_0 = 2$, $a_{i+1} = a_i + r(a_i)$ when $x < 3/2$ and $a_{i+1} = a_i - r(a_i)$ when $x > 3/2$. This will yield an infinite sequence unless we at some point hit $3/2$, which would require a very fine-tuned r .

A more promising idea is to rule out models whose domains are in some sense “too small”. In \mathfrak{M} , we can get from any point to any other in at most three accessibility steps, and even in the n -dimensional variants discussed above, we can still get from any point to any other in at most five steps. In reality, on the other hand, there are surely possible interpretations that can be reached from the correct interpretation only by *many* steps (where each step involves passing from an interpretation to another that stands to it in the relation to which it maps ‘indiscriminable’). But we can easily construct models that with cutoffs at every order in which there are points that are arbitrarily many steps apart, just by putting countably many copies of \mathfrak{M} “back-to-back”.¹³

So, finding well-motivated, formally tractable constraints which guarantee a maximum cutoff will not be easy. But this is not yet an argument that people could in fact speak a language in which some sentence was vague at order n and precise at higher orders. To address that philosophical question with the methods of model theory, one would need to consider what “realistic” models of the vagueness of speakable languages would look like. And there are good reasons to think that none of the models we have considered so far are very realistic. In particular, their domains do not contain points with tiny or enormous accessibility radii (corresponding to extremely demanding or undemanding interpretations of ‘indiscriminable’), let alone points corresponding to interpretations on which ‘indiscriminable’ expresses a completely different kind of relation altogether. The absence of such very distant points is problematic from the point of view of realism, since it is plausible that even interpretations radically dissimilar from the correct interpretation can be reached from it

¹³To be precise: let the domain be the real line, with the standard metric. Where $\lfloor x \rfloor$ is the greatest integer $\leq x$, let $r(x) = (1 + x - \lfloor x \rfloor)/3$ if $\lfloor x \rfloor$ is odd and $r(x) = (2 + \lfloor x \rfloor - x)/3$ if $\lfloor x \rfloor$ is even. Let $\mathbb{I} = \bigcup_i (2 + z - a_i, z + a_i)$ for all even integers z .

in finitely many steps, so that they will eventually need to be considered as we evaluate sentences with more and more iterations of ‘ Δ ’. However, once we start taking this thought seriously, it becomes hard to see how *any sentence at all*—even a logical truth—could remain true when prefaced by arbitrarily many ‘definitely’s. After all, we can get in finitely many steps of arbitrarily small size from the actual world to a world where the use of the logical constants is completely different. I have attempted elsewhere [citation deleted for blind review] to turn these worries into an argument that every sentence is vague at arbitrarily high orders. Those who disagree with this conclusion will need to find some way to sustain the thought that some interpretations are so distant from the correct one that they should be excluded altogether: this will make it hard to argue that our models must be unrealistic on the grounds that their domains are insufficiently varied. Those who agree with me, on the other hand, will have to give up the assumption that ‘definitely’ has a normal modal logic, and with it the assumption that any normal modal model is perfectly “realistic” in the relevant sense. This perspective requires a different perspective on the philosophical significance of such models. They may still be useful tools for probing certain questions about vagueness, but we will need to take some of their features with a grain of salt, particularly features that turn on deep embeddings of Δ , or that depend on whether or not the model includes points representing interpretations that are very far from being correct. I would still like to know whether the possibility of high cutoffs is in some sense a “typical” feature of variable radius frames, or whether it always requires the kind of fine-tuning that characterises the examples we have found. But I see no easy route from an answer to this question to any conclusion about how higher-order vagueness actually works.

References

- Bacon, Andrew (MS). ‘Vagueness at Every Order. The Prospects of Denying B’.
- Caie, Michael (2012). ‘Vagueness and Semantic Indiscriminability’. *Philosophical Studies* 160, pp. 365–77.

- Graff, Delia (2002). 'An Anti-Epistemicist Consequence of Margin for Error Semantics for Knowledge'. *Philosophy and Phenomenological Research* 64, 127—142.
- Kearns, Stephen and Ofra Magidor (2008). 'Epistemicism about Vagueness and Metalinguistic Safety'. *Philosophical Perspectives* 22, pp. 277–304.
- Mahtani, Anna (2008). 'Can Vagueness Cut Out at Any Order?' *Australasian Journal of Philosophy* 86, pp. 499–508.
- Williamson, Timothy (1990). *Identity and Discrimination*. Oxford: Blackwell.
- (1994). *Vagueness*. London: Routledge.
- (1999). 'On The Structure of Higher-Order Vagueness'. *Mind* 108, pp. 127–43.
- (2000). *Knowledge and Its Limits*. Oxford: Oxford University Press.