

Iterating Definiteness

Cian Dorr

Final version: December 1, 2008

1 Preliminaries

A central concept in the study of vagueness is the concept of a borderline case. This concept has its most basic application when we are faced with a question of the form ‘Is x F ?’, but are unwilling to answer ‘Yes’ or ‘No’ for a certain distinctive kind of reason. Wanting to be co-operative, we need to say *something*; by saying ‘It’s a borderline case’, we excuse our failure to give a straightforward answer while conveying some information likely to be of interest to the questioner. The nature of the considerations that make us unwilling to answer ‘Yes’ or ‘No’ in these cases is a topic of central importance in the philosophy of vagueness. Different views about this naturally lead to different answers to the question what it means to be a borderline case.

Before we can even broach this question, we need to settle on a way of regimenting borderlineness-talk. There are two main approaches. On the first approach, the basic notion is metalinguistic in character, so that the task is that of making sense of locutions like these:

Sentence S is borderline as used by community C .

Sentence S is borderline as used by community C at possible world w .

$\langle x_1, \dots, x_n \rangle$ is a borderline case of predicate Φ as used in contexts of type T by community C at possible world w .

On this approach, the philosophy of vagueness is clearly a branch of the philosophy of language. This is less clear on the second approach, which regiments ‘borderline’ as an *operator* (see Fine 1975, p. 148ff.). On this view, ‘It is borderline whether P ’ is no more a claim about language than ‘It is contingent whether P ’. The fact that it is borderline whether P , if it is a fact, is not a fact especially about any particular community or any particular linguistic expression; it can be expressed equally straightforwardly in many different languages.¹

The dispute is about priority: there is no reason for partisans of either approach to reject the vocabulary of the other approach as unintelligible. Suppose we already understand ‘it is borderline whether’ and the notion of a sentence being *true* as used by a community at a world. Then we can analyse ‘ S is borderline as used by C at w ’ as ‘It is borderline whether S is true as used by C at w ’ (cf. Fine 1975, p. 296). Conversely, if we already understand ‘ S is borderline as used by C at w ’, we could define the borderlineness operator by stipulating that ‘it is borderline whether Φ ’ should be synonymous with ‘ Φ is borderline as used by us at the actual world’ (when Φ is a closed sentence) or ‘ $\langle v_1, \dots, v_n \rangle$ is a borderline case of ‘ Φ ’ as used by us at the actual world’ (when Φ is an open sentence with free variables $v_1 \dots v_n$).^{2,3}

¹Perhaps we should think of ‘it is borderline whether P ’ as ascribing a property, borderlineness, to a proposition, the proposition that P , or some kindred abstract, non-linguistic entity. Or perhaps we should resist such attempts to impose a subject-predicate structure on sentences constructed using operators, as Prior did for modal and temporal operators (Prior 1968). This is an interesting subsidiary dispute, but as far as our overall conception of the place of vagueness in the scheme of things is concerned, the dispute between the metalinguistic and non-metalinguistic approaches is more central.

²I use Quine’s corner-quotes (Quine 1940). Since ‘ Φ ’ is a variable ranging over linguistic expressions, ‘‘ Φ is borderline as used by us at the actual world’’ is synonymous with ‘the result of writing “” and then Φ and then “” is borderline as used by us at the actual world’’’.

³Opponents of the metalinguistic approach will complain that the operator defined in this way doesn’t interact in the right way with ambiguous expressions: ‘it is borderline whether S ’ should be ambiguous whenever S is, but if ‘borderline’ and quote-names are not ambiguous, ‘‘ S is borderline as used by us at the actual world’’ will never be ambiguous. One could attempt to fix this by treating ambiguity as homonymy, so that, e.g., the quote-name ‘‘Some banks are closed’’ is ambiguous, referring on two different disambiguations to different linguistic entities. But it is not clear how this could work. If we are to understand borderlineness as a feature of a sentence’s use, we had better take sentences to be entities that can be used in different ways, with different meanings, by different communities at different possible worlds. This makes it hard to give an account of what the difference between the putatively distinct items named by ‘Some banks are

It is customary and convenient to treat ‘it is borderline whether’ as defined in terms of a ‘definitely’ operator, with ‘it is borderline whether P ’ analysed as ‘Not definitely P and not definitely not- P ’—in symbols, $\neg\Delta P \wedge \neg\Delta\neg P$. We could either treat this as a first step in the analysis of ‘it is borderline whether’, or—if we prefer doing things the other way round—as an elementary logical consequence of the analysis of ‘definitely P ’ as ‘ P and it is not borderline whether P ’.

On the metalinguistic approach, it is similarly traditional to analyse borderlineness in terms of truth and falsehood: ‘ S is borderline as used by C at w ’ is analysed as ‘ S is neither true nor false as used by C at w ’. This is a controversial move: it is not so clear that our intuitive notions of truth and falsehood for sentences behave as they would need to behave for this analysis to be tenable (Williamson 1994, section 7.2). But those who are wary of this analysis should at least agree that there are two distinctively different ways in which a meaningful sentence can fail to be borderline, even if ‘true’ and ‘false’ aren’t the right labels for these ways. To preserve neutrality I’ll speak of ‘dtrue’ and ‘dfalsehood’. If you like, you can pronounce ‘dtrue’ as ‘definitely true’, and ‘dfalse’ as ‘definitely false’; but of course only followers of the operator approach will want to analyse these notions as the result of applying the ‘definitely’ operator to antecedently understood notions of truth and falsity. Followers of the metalinguistic approach should instead—at least as a first approximation—analyse ‘Definitely S ’ as ‘‘ S ’ is dtrue as used by us at the actual world’.

Both conceptions of the relation between ‘definitely’ and ‘dtrue’ vindicate the following principle, which we will need to refer back to later:

T- Δ For any sentence S , ‘ ΔS ’ is dtrue as used by us at the actual world iff definitely, S is dtrue as used by us at the actual world.

On the operator approach, analysing ‘dtrue’ as ‘definitely true’, instances of T- Δ can be derived closed’ could be, or of what would make it the case that one rather than the other of them was used by a given community. Moreover, adherents of the metalinguistic approach are liable to hold that ambiguity and vagueness are kindred phenomena, which should be treated in a unified fashion at the most fundamental level of theorising. (For a radical version of this, see Braun & Sider 2007.) If so, even if we can make sense of a notion of sentences as uninterpreted but disambiguated, they will be entities of at best secondary theoretical importance. Nevertheless, it may be possible to find a way of talking about them that is adequate to the task of analysing ‘it is borderline whether’ so that it conforms to the rule that ‘it is borderline whether S ’ inherits any ambiguity in S .

from instances of the T -schema strengthened by two ‘definitely’ operators: ‘Definitely, definitely, ‘ S ’ is true as used by us at the actual world iff S ’.⁴ On the metalinguistic approach, we can argue as follows, appealing to the analysis of ‘ $\Delta\Phi(v)$ ’ as ‘ Φ ’ is dtrue of v as used by us at the actual world’:

- (1) For any sentence S , ‘ ΔS ’ is dtrue as used by us at the actual world iff ‘ S ’ is dtrue as used by us at the actual world’ is dtrue as used by us at the actual world.
- (2) For any expression Φ and unary predicate F , ‘ Φ ’ is F ’ is dtrue as used by us at the actual world iff F is dtrue of Φ as used by us at the actual world.
- (3) So for any S , ‘ ΔS ’ is dtrue as used by us at the actual world iff ‘dtrue as used by us at the actual world’ is dtrue of S as used by us at the actual world.
- (4) So for any S , ‘ ΔS ’ is dtrue as used by us at the actual world iff definitely, S is dtrue as used by us at the actual world.⁵

2 Infinite definiteness

Once we have introduced the ‘definitely’ operator in one way or another, it becomes natural to think about stronger operators defined by iterating it. We have the sequence $\Delta, \Delta\Delta, \dots, \Delta^i, \dots$. And there are various ways we can introduce something like an infinite limit to that sequence. The most straightforward is to use infinitary conjunction, defining ‘ $\Delta^\omega S$ ’ as ‘ $S \wedge \Delta S \wedge \Delta\Delta S \wedge \dots \wedge \Delta^i S \wedge \dots$ ’. In my view, this is legitimate: while English is not itself an infinitary language, the ellipsis ‘...’

⁴If the semantic paradoxes force us to reject the claim that all instances of the T -schema are true, this won’t be enough for a general argument for $T\Delta$. But since the semantic paradoxes don’t undermine $T\Delta$ itself, followers of the operator approach will presumably want to hold onto it in any case, just as with analogous Tarski-style principles for other operators.

⁵This argument should still go through even if we reject the proposed analysis of ‘Definitely S ’ in favour of something subtler, in order to accommodate the claim that ‘Definitely S ’ inherits the ambiguity of S . For since ‘dtrue’ is not itself ambiguous, (1) will remain plausible, and the inference from (3) to (4) will remain valid, whether we understand the notion of dtruth applied to perhaps-ambiguous sentences as requiring dtruth on all disambiguations or merely on some.

lets us express in English some definitions that would otherwise be formulable only in an infinitary language.

Those who are suspicious of infinitary conjunctions can use quantification to achieve more or less the same effect. Say that S is *ultratrue as used by C at w* iff every finite definitization of S is dtrue as used by C at w , where a *finite definitization* of S is a sentence that consists of S preceded by zero or more Δ s. As an alternative to an infinitary conjunction, we could define $\ulcorner \Delta^\omega S \urcorner$ as $\ulcorner S \urcorner$ is ultratrue as used by us in the actual world⁶.

On either definition, the operator Δ^ω is a puzzling one. One source of puzzlement is the apparent validity of the following schema:

$$\text{DEF} \quad \Delta^\omega P \rightarrow \Delta \Delta^\omega P.$$

If we define Δ^ω using an infinite conjunction, each instance of DEF will follow from DIST, which is the obvious extension to the infinitary case of the principle that definiteness distributes over conjunction:

$$\text{DIST} \quad \Delta(P_1 \wedge P_2 \wedge \dots) \leftrightarrow (\Delta P_1 \wedge \Delta P_2 \wedge \dots)^7$$

If on the other hand we define $\ulcorner \Delta^\omega S \urcorner$ as $\ulcorner S \urcorner$ is ultratrue as used by us at the actual world⁷, we can argue for DEF by appealing to T- Δ :

- (1) Every finite definitization of S is dtrue as used by us at the actual world (premise).
- (2) T is a finite definitization of S (assumption).
- (3) $\ulcorner \Delta T \urcorner$ is a finite definitization of S (definition of ‘finite definitization’).
- (4) $\ulcorner \Delta T \urcorner$ is dtrue as used by us at the actual world (1, 3).
- (5) Definitely, T is dtrue as used by us at the actual world (4, T- Δ).

⁶We could equally well have used ‘true’ instead of ‘dtrue’ in this definition; an argument similar to the one below shows that the definition with ‘true’ entails the one with ‘dtrue’.

⁷Use conjunction elimination to get from $P \wedge \Delta P \wedge \Delta \Delta P \wedge \dots$ to $\Delta P \wedge \Delta \Delta P \wedge \dots$, then DIST to get $\ulcorner \Delta(P \wedge \Delta P \wedge \Delta \Delta P \wedge \dots) \urcorner$.

- (6) Every finite definitization of S is definitely dtrue as used by us at the actual world (2–5).
- (7) If there are some things which are definitely all and only the F s, then every F is definitely a G iff definitely, every F is a G (premise schema).⁸
- (8) There are some things that are definitely all and only the finite definitizations of S (premise).⁹
- (9) Definitely, every finite definitization of S is dtrue as used by us at the actual world (6, 7, 8).

The key premise here is (7), which does the same work in this argument that DIST did the previous argument, and seems similarly plausible.¹⁰

DEF makes a certain kind of trouble for someone who thinks that ultratruth is a common phenomenon. Let's say that x is *ultrabald* iff $\Delta^\omega(x \text{ is bald})$. By DEF, whenever one is ultrabald, one is definitely ultrabald. So whenever it is borderline whether one is ultrabald, one is not ultrabald.

This is not yet to say that 'ultrabald' is precise, or that no-one is borderline ultrabald. To show that, we would also need an argument that everyone who is *not* ultrabald is definitely not

⁸(7) is a consequence of the principle that definiteness commutes with universal quantification: $\Delta\forall x(\Phi(x)) \leftrightarrow \forall x(\Delta\Phi(x))$. But unlike that principle, (7) leaves room for vagueness deriving from the quantifiers as well as vagueness derived from the predicates.

⁹I assume that any vagueness in the reference of quote-names—which abstract entity, precisely, does 'Δ' refer to?—can be harmlessly ignored.

¹⁰I don't mean to suggest that DIST or (7) is beyond dispute. Hartry Field (2003b; 2008) has recently argued for rejecting such principles, on the grounds that doing so makes available a resolution of the semantic paradoxes that preserves the full intersubstitutability of ϕ and $\ulcorner \text{True}(\langle\phi\rangle) \urcorner$ (where $\ulcorner \langle\phi\rangle \urcorner$ denotes the Gödel number of ϕ) while validating the inferences (i) $P \models \Delta P$ and (ii) $P \rightarrow \neg P \models \neg \Delta P$. Where Q_ω is $\neg \Delta^\omega \text{True}(\langle Q_\omega \rangle)$, we have $\text{True}(\langle Q_\omega \rangle) \rightarrow \neg \Delta^\omega \text{True}(\langle Q_\omega \rangle)$ by intersubstitutivity, and hence $\Delta^\omega \text{True}(\langle Q_\omega \rangle) \rightarrow \neg \Delta^\omega \text{True}(\langle Q_\omega \rangle)$ by the factivity of Δ^ω . It follows by (ii) that $\neg \Delta \Delta^\omega \text{True}(\langle Q_\omega \rangle)$. Given DEF, we could infer from this that $\neg \Delta^\omega \text{True}(\langle Q_\omega \rangle)$, i.e. to Q_ω itself. Then intersubstitutivity would give us $\text{True}(\langle Q_\omega \rangle)$, repeated applications of (i) would yield $\Delta^k \text{True}(\langle Q_\omega \rangle)$ for each finite k , and finally, by an infinite conjunction introduction, we would have $\Delta^\omega \text{True}(\langle Q_\omega \rangle)$: a contradiction. It is not clear to me to what extent someone who accepts Field's view should expect DEF to fail even for ordinary non-semantic vague predicates. Incidentally, one moral of the argument I will be giving below is that (i) is *not* valid.

ultrabald.¹¹ But it does mean that we won't be able to use the notion of a borderline case in the usual way to excuse our failure to give straightforward 'Yes' or 'No' answers to questions about ultrabaldness. If I'm asked 'Is so-and-so ultrabald?' and for some reason I don't want to commit myself to the extent of saying 'Yes' or 'No', I should be just as unwilling to say 'He's a borderline case'; if I say this, I will have asserted something at least as strong as what I would have asserted by saying 'No'.

You will not be embarrassed by this if you think you know some precise necessary and sufficient condition for ultrabaldness. Otherwise, you may find it hard to respond co-operatively to questions about ultrabaldness without being able to appeal to borderlineness in the usual way. Inevitably, there will be cases where you will be unwilling to answer 'Yes' or 'No' to the question 'Is this person ultrabald?', no matter how much you might learn about the relevant precise facts. What should you say, given that saying 'Borderline' would commit you to saying 'No'? Should you answer randomly? Should you just remain silent? These options are hardly consistent with the standard of co-operativeness to which you are trying to hold yourself.

You will be tempted to say 'I don't know'. This is, after all, what we standardly say when we want to be co-operative but don't want to give a straightforward answer to a question. But there are various reasons why we might be uncomfortable with such a response. In other work (Dorr 2003) I have argued that in many ordinary cases where it is borderline whether *P*, and one is reasonably well informed about the relevant underlying facts, it is borderline whether one knows that *P*. On this view, unless we can identify some special reason why knowledge would be harder to come by in borderline cases of ultrabaldness than in other borderline cases, 'I don't know' is liable to be just as unacceptable an answer as 'Yes' or 'No'.

But even if you hold the more orthodox view that knowledge is inconsistent with borderlineness, you may still find there to be something unsatisfying about simply admitting that we don't know who is and is not ultrabald and leaving it at that. Shouldn't we philosophers who take a

¹¹The most obvious route to that claim would involve appealing to the controversial B (Brouweresche) axiom schema $\neg P \rightarrow \Delta \neg \Delta P$, or perhaps to some weaker axiom schema of the form $\neg P \rightarrow \Delta \neg \Delta^n P$ —see Graff 2002.

professional interest in questions of vagueness want to know more? If you have admitted that you don't know whether someone is ultrabald, despite having been given as much time to reflect and as much access to other relevant facts as you have any use for, you will probably react with impatience to the suggestion that you undertake further inquiries. You will be tempted to protest that such inquiries would be pointless: you don't just *happen* not to know; rather, neither you nor anyone else is even *in a position* to know, given any amount of further inquiry.

But what could explain this inability? If borderlineness is a barrier to knowledge, your inability to know whether x is ultrabald *might* be explained by its being a borderline case. But since you don't know that x is not ultrabald, you don't know that it is borderline whether x is ultrabald, so you must leave open the possibility that the obstacle to your knowing whether x is ultrabald is of some other kind. But what other sort of obstacle to knowledge could be relevant in this context? Whatever it is, why can't we tell it apart from the obstacle to knowledge characteristic of borderlineness? Wouldn't it make more sense to adopt a more expansive use of the expression 'borderline', on which it applies to cases in which either sort of obstacle is present? If we did adopt this more expansive sense of 'borderline', along with corresponding senses of 'definitely', 'dtrue', 'ultratrue' and 'ultrabald', what then would be our epistemic situation with respect to the question 'what does it take to be ultrabald?'

If there are people who are borderline ultrabald, and borderlineness is the only relevant obstacle to knowledge, we are doomed never to achieve a certain kind of theoretical satisfaction in our relations with them. So long as we form no opinion on whether they are ultrabald, we will never know whether, in failing to form an opinion, we are passing up knowledge which there is no obstacle to our possessing. If we do in fact bring give up on further inquiry close, we will always be wondering if we could have resolved the question just by giving it a bit more thought.

This would be an unsettling conclusion, I think. But these considerations don't rise to the level of an argument that 'ultrabald' has precise and knowable conditions of application; at best, they show why it would be nice to have such an argument. The task of giving one will occupy the remainder of the paper. My conclusion will be radical: no-one is ultrabald; in faact no sentence

whatsoever is ultratrue, and no predicate whatsoever is ultratrue of anything.

3 No sentence is ultratrue: first attempt

In this section I will take a first stab at arguing that no sentence is ultratrue. The result won't be terribly hard to resist, but will serve as a basis for later refinements.

It will help to make some simplifying assumptions about the laws of nature that enable us to define a well-behaved notion of distance between nomologically possible worlds. Suppose, then, that the actual world consists of finitely many point-particles in Newtonian absolute space. Where w and w' are nomologically possible worlds with the same particles, define the distance between w and w' at t as the sum, for each particle, of the distance between the point where the particle is located at t at w and the point where it is located at t at w' .¹² Assume too that the laws of nature are deterministic and continuous, so that for any δ and t , there is a δ' such that any two worlds which are less than δ' apart now will remain less than δ apart until at least t units of time hence.

With some such notion of inter-world distance in hand, we can state our argument. It has two premises:

SERIES For every positive real number δ and sentence S , there is a sequence $\langle w_0, \dots, w_n \rangle$ of possible worlds such that:

- S1 S is not dtrue as used by us at w_0 .
- S2 w_n is the actual world.
- S3 For each $0 \leq i < n$, the distance between w_i and w_{i+1} is less than δ .
- S4 Our use of ' Δ ' at each w_i is at least as stringent as it is at the actual world, in the following sense: for any sentence T , necessarily, if $\ulcorner \Delta T \urcorner$ is dtrue as used by us at w_i , then definitely, T is dtrue as used by us at w_i .

¹²We could equally well use the notion of distance in configuration space standard in physics, which brings in the masses of the particles. Or, for an alternative that doesn't require identifying points of space across possible worlds, see Barbour 2006.

MARGIN There is a $\delta > 0$ such that whenever a sentence S is definitely dtrue as used by us at w , and the distance between w and w' is less than δ , S is dtrue as used by us at w' .

Now for the argument. Let δ meet the condition specified by MARGIN; let S be an arbitrary sentence; let $\langle w_0, \dots, w_n \rangle$ be a sequence satisfying S1–S4. We show by induction that for each $m \leq n$, $\lceil \Delta^m S \rceil$ is not dtrue as used by us at w_m . The base step is just S1. For the induction step, assume that some sentence T is not dtrue as used by us at w_m . By S3 and our choice of δ , T is not definitely dtrue as used by us at w_{m+1} ; and so (by the contrapositive of S4), ΔT is not dtrue as used by us at w_{m+1} . So in particular, if $\lceil \Delta^m S \rceil$ is not dtrue as used by us at w_m , $\lceil \Delta^{m+1} S \rceil$ is not dtrue as used by us at w_{m+1} , which is what we need for the induction. Letting $n = m$, then, we have that $\lceil \Delta^n S \rceil$ is not dtrue as used by us at w_n , i.e. at the actual world. A fortiori, S is not ultratrue, and $\lceil \Delta^\omega S \rceil$ is not dtrue, as used by us at the actual world.

Some comments, before we discuss how the premises SERIES and MARGIN might be justified:

- (a) At this point, it will be best to interpret predications of the form ‘ S is dtrue as used by C at w ’ as having to do with dtruth *at the actual world* rather than dtruth *at w* .¹³ If we interpreted such predications as having to do with the dtruth *at w* of sentences as used at w , MARGIN would be implausible, since it would rule out our ever introducing perfectly precise sentences which express nomologically contingent truths. Framing such precise distinctions is difficult but not impossible. The expression ‘one second’, currently defined as ‘the duration of 9,192,631,770 periods of the radiation corresponding to the transition between the two hyperfine levels of the ground state of the caesium-133 atom’ (BIPM), is a reasonable candidate for being perfectly precise.¹⁴ If it is, and the duration of the universe is finite, there will be sentences of the form ‘The universe lasts for at least n seconds’ which are definitely dtrue at the actual world and definitely dfalse at worlds arbitrarily close to the actual world, not only as used at the actual world, but also as used at those worlds.

¹³I will reconsider this choice later, in section 5.

¹⁴‘One Planck time’ is an even better candidate.

- (b) Worries about the applicability of mathematical induction to vague predicates are not really to the point. It should be straightforward to argue for some large finite bound on the lengths of the sequences we need to consider, in which case we could reconstruct the argument using finitely many applications of modus ponens.
- (c) We can run an exactly similar argument for the claim that no predicate is ultratrue of any sequence of arguments. But for simplicity I will continue to focus when possible on closed sentences.
- (d) I have suppressed all mention of time. *A*-theorists about time shouldn't mind this. *B*-theorists should either take every sentence as implicitly relativised to the present time, or else reinterpret all claims about "possible worlds" as claims about ordered pairs of worlds and times.
- (e) I have suppressed the apparatus that would be necessary to deal with context-sensitive sentences. It should be easy to reintroduce.
- (f) Since sentences need to be understood as items which can be used very differently at different worlds, they are presumably the sort of thing that can be ambiguous (as used by a given community at a given world). There are two ways of thinking about what it might mean to describe such entities as 'dtrue': it could mean 'dtrue on every disambiguation' or 'dtrue on some disambiguation'. As far as I can see, it makes no difference which of these we adopt.

The case for SERIES is straightforward. Evidently we *could* have used any sentence *S* in such a way as to make it not be dtrue at the actual world, e.g. by using *S* in the same way that we actually use the sentence ' $0 = 1$ ', or by speaking a language in which *S* is not meaningful at all. We could have done this even at a world with the same kinds of laws we have been supposing to hold at the actual world, and with the same number of particles as the actual world. But any such world is a finite distance from the actual world, and thus can be reached from the actual world by way of

finitely many steps of arbitrarily small size. The only remaining question is whether we can take these steps in such a way that all the worlds we visit along the way satisfy S4. By T- Δ , the actual world satisfies S4. There is no trouble choosing w_0 in such a way as to satisfy S4: for example, we could let w_0 be a world where ‘ Δ ’ is not meaningful at all, or where it is used so demandingly that no sentence of the form ‘ ΔS ’ is dtrue. And I see no reason to doubt that at least some such w_0 can be connected to the actual world by a sequence of worlds all satisfying S4. If we think of ‘definitely’ as expressing a property of propositions, it would suffice for the property it expresses at each w_i where it is meaningful at all to be the same as, or stronger than, the one it expresses at the actual world. If we think of ‘ ΔS ’ as meaning something like ‘‘ S ’ is dtrue as used by us at the actual world’’, it would suffice for this equivalence to be in place at each w_i , and the use of ‘dtrue’ at each w_i to be such that necessarily, whenever ‘dtrue as used by... at ...’ is dtrue of some $\langle S, C, w \rangle$ as used by us at w_i , it is dtrue of $\langle S, C, w \rangle$ as used by us at the actual world. In either case, it is hard to see how their could fail to be a topologically connected set of worlds satisfying S4 which contains both the actual world and some appropriate w_0 .

So much for SERIES. Why would anyone accept MARGIN?

If you are anything like me, you will have noticed an affinity between MARGIN and certain claims characteristic of Timothy Williamson’s epistemic theory of vagueness (1994). In fact, I think I see a good argument from Williamson’s view to MARGIN. But the case is less straightforward than I initially supposed.

For Williamson, the claim that it is definitely the case that P means, or at least entails, that there is no obstacle of a certain distinctive kind to our knowing that P . It is sufficient for the existence of such an obstacle for there to be a false proposition which we could very easily have used the sentence we actually use to express the proposition that P to express. That is:

W1 If we use S and no other sentence to express proposition p at w and to express proposition q at a world w' that is close to w , then at w : if p is definitely true, q is true.

If we could drop the ‘at w ’ from W1, we’d have something from which we could hope to derive MARGIN, at least restricted to communities in which each proposition is expressed by at most one

sentence. It would just be a matter of putting a ‘definitely’ in front of W1, and arguing that for some δ , whenever the distance between two possible worlds is less than δ , they are definitely “close” in the relevant sense. But what could license eliminating the ‘at w ’? We might try arguing as follows. First, insert actuality operators in the consequent of W1, so that it becomes

W2 ... at w : if definitely (p is actually true), then (q is actually true).

Next move ‘at w ’ inside the conditional:

W3 ... if at w , definitely (p is actually true), then at w , (q is actually true).

Then interchange ‘at w ’ and ‘definitely’:

W4 ... if definitely (at w , p is actually true), then at w , (q is actually true).

Finally, appeal to the definite validity of the logical schema ‘ P iff at w actually Φ ’, to get

W5 ... if definitely (p is true), then q is true

which is what we wanted.

Unfortunately, two of these steps are dubious by Williamson’s lights. The step from W1 to W2 is problematic chiefly because it is hard to understand the question whether there is an obstacle of the relevant kind at w to our knowing that p is *actually* true. How are people at w supposed to pick out the actual world at all, in order to formulate the question whether p is true at it? On the most straightforward way of understanding what this would require (Williamson 1987; Soames 1998), the obstacles at w to our even *entertaining* p is actually true are so formidable that it is hard even to make sense of the question whether there are, in addition, any obstacles of the distinctive sort associated with vagueness to our knowing it.

The second dubious step is from W3 to W4. While it is tempting in reasoning about vagueness to treat ‘definitely’ as commuting with ‘at w ’, there is no obvious support for this in Williamson’s theory. In general, the claim that there is no obstacle of some given sort to our knowing that at w , P is independent of the claim that at w , there is no obstacle of that sort to our knowing that P . And

there is no obvious reason why the particular sort of obstacle to knowledge that is distinctive of vagueness on Williamson's view should be exceptional in this regard.¹⁵

Thus, the most obvious route from Williamson's theory to MARGIN is fraught with difficulties. But we can do better, by focusing not on the use of the sentence S at w and w' , but at the use of the predicate 'dtrue as used by ... at ...' at worlds close to the actual world.

This predicate is manifestly vague. For any S , it is easy to construct Sorites sequences of possible worlds which take us in many small steps from a world of which it is clearly the case that S is dtrue as used by us there to a world of which this is clearly not the case. These sequences raise the same puzzles as the canonical Sorites sequences involving 'bald' and 'heap'.¹⁶ While we can see that the accumulated tiny differences must somehow constitute the difference between a way of using S that makes it dtrue (at the actual world) and one that doesn't, we have no more grip on the question how any one step along the sequence could constitute such a difference than we have on the question how a similarly tiny difference between two worlds could make it be the case that I am bald at one world and not at the other. We thus have as much reason to recognise that it is sometimes borderline whether S is dtrue as used by us at w as we do to recognise that it is sometimes borderline whether I am bald at w .¹⁷ Moreover, just as we have reason to believe

¹⁵The failure of 'definitely' to commute with 'at w ' on Williamson's view makes it surprisingly hard for propositions to be necessarily definitely true. For it to be false that necessarily definitely $0 \neq 1$, it would suffice for there to be a pair of close worlds w, w' , such that a sentence S that is used at w to express the proposition that $0 \neq 1$ is used at w' to express something false at w . But given that it must be possible for a sentence to express the proposition that $0 \neq 1$ even though it could easily have expressed *some* other proposition, there is no obvious reason why it shouldn't be possible for a sentence to express the proposition that $0 \neq 1$ even when it could easily have expressed a *false* proposition.

¹⁶Cf. Sorensen 1985.

¹⁷John Hawthorne (2006) notices special problems with the idea that 'true' (or 'true as used by ... at ...') expresses different things at nearby worlds. Unless instances of the disquotation schema are in danger of expressing falsehoods, each relation R that is expressed by 'true as used by ... at ...' at some w near the actual world must be such that, necessarily, for any S , $R(S, \text{us}, w)$ iff S is true as used by us at w . So there won't be worlds where 'true as used by ... at ...' expresses relations that are *uniformly* more demanding than the relation expressed by this predicate at the actual world. But the claim in the text concerns 'dtrue' rather than 'true'. Even if the intension of 'true' were the same at all nearby worlds, we could still get the intension of 'dtrue' to vary by varying the use of 'definitely'.

that it will never happen, for sufficiently close worlds w and w' , that I am definitely bald at w and definitely not bald at w' , we have reason to believe that it will never happen S is definitely dtrue as used by us at w and definitely not dtrue as used by us at w' . That is:

BORDERLINE There is some $\delta > 0$ such that whenever a sentence S is definitely dtrue as used by us at w , and the distance between w' and w is less than δ , S is not definitely not dtrue as used by us at w' .

To help make this plausible, consider a very tiny value of δ , such that given the laws, whenever the distance between w and w' is less than δ , it will take a trillion years before there is as much as a nanometre's difference between the location of any particle at w and its location at w' . Our conception of the way in which dtruth-conditions depend on use seems far too imprecise for such a tiny difference in use ever to make the difference between definite dtruth and definite lack of dtruth (at the actual world, or indeed at any given world).¹⁸

Since 'not definitely not dtrue' is weaker than 'dtrue', **BORDERLINE** is, formally speaking, weaker than **MARGIN**. But Williamson's theory of vagueness provides a way of closing the gap. For Williamson, the vagueness of any predicate consists in the fact that there are worlds close to the actual world where it expresses relations different in intension from the one it actually expresses. In principle, the actual world could be a "local maximum" with respect to the use of some predicate, so that the relation it actually expressed entailed all the relations it expressed at nearby worlds. But on any remotely plausible account of the connection between use and meaning, this will happen only in very special cases. Normally, if there are close worlds where a predicate expresses a relation weaker than the one it actually expresses, say because its use is slightly laxer in some respects, there will also be close worlds where it expresses a relation stronger than the one it actually expresses, because its use slightly stricter in those same respects.¹⁹

¹⁸But see section 5 below for an important objection to this claim.

¹⁹Plausibly, for this to fail, the intension actually expressed would have to be fairly "natural", since a range of different patterns of use including the actual pattern together with all the "stricter" patterns that obtain in close worlds result in the same intension being expressed. But the pull of naturalness cannot be too strong: the actual world must be perched near the edge of the set of

I see no reason to think that ‘dtrue’ would be abnormal in this respect. In fact, on Williamson’s theory, it there will plausibly be worlds close to the actual world where ‘dtrue as used by... at...’ expresses a relation that is uniformly stronger than the one it actually expresses, in the sense that whenever the former relation holds between S , C and w , the latter holds between S , C , and any w' within δ of w , for some δ . One would expect this to happen if people at the close world in question are, across the board, a bit more reluctant to apply the predicate ‘knows’ than they actually are, with the result that ‘knows’ expresses a uniformly stronger relation R , which a person bears to a proposition at w only if the person knows the proposition at all w' sufficiently close to w . It is thus easier for there to be obstacles to the obtaining of R than for there to be obstacles to knowledge. If the connections Williamson posits between ‘knows’, ‘definitely’, ‘true’ and ‘dtrue’ remain in place at the world in question, this will make it harder for a proposition to have the property expressed there by ‘definitely true’, and thus harder for a sentence to stand in the relation expressed there by ‘dtrue’ to any given community and world.

To get from BORDERLINE to MARGIN, we only need a weaker claim: roughly, that if there are close worlds where the relation expressed by ‘dtrue’ is weaker along some dimension, there are close worlds where it is stronger along that dimension. To make this precise, let a “dimension” be a triple $\langle S, C, \lambda \rangle$, where λ is a straight path through the space of nomologically possible worlds, which starts with a world w^+ such that S is definitely dtrue as used by C at w^+ , ends at a w^- such that S is definitely not dtrue as used by C at w^- , and is such that for every other $w \in \lambda$, it is borderline whether S is dtrue as used by C at w . The actual world is a local maximum with respect to this dimension if the relation expressed by ‘dtrue’ at the actual world is one that fails to hold between S , C and w for any $w \in \lambda$ other than w^+ . Given that there are no especially natural relations in the vicinity of ‘dtrue’ is especially natural, it is plausible that this never happens. A stronger claim also seems plausible: that the actual world does not come *arbitrarily close* to being a local maximum—in other words, for some n , the actual cutoff for ‘dtrue’ occurs at least $1/n$ of the way along each $\langle S, C, \lambda \rangle$, measured by our canonical notion of interworld distance. This gives

worlds in which this intension is expressed, since different intensions are expressed at close worlds with “laxer” patterns of use where different intensions are expressed.

us what we need: since by BORDERLINE, there is a δ such that the length of λ is always at least δ , it follows that the distance along λ between w^+ and any world w such that S is not dtrue as used by C at w is always at least δ/n . Thus δ/n witnesses the truth of MARGIN: whenever a sentence S is definitely dtrue as used by C at w , and the distance between w' and w is less than δ/n , S is dtrue as used by C at w' .

The conclusion that no sentence is ultratrue need not be unwelcome or even especially surprising to Williamson. A central doctrine of Williamson 2000 is that only trivial conditions are *luminous*: such that necessarily, whenever they obtain, one is in a position to know that they obtain. This strongly suggests that the claim that P entails that ΔP only when it is trivial that P . Since ' $\Delta^\omega P \vdash \Delta\Delta^\omega P$ ' is valid, and since it would have to be trivial that P for it to be trivial that $\Delta^\omega P$, it follows that that all ultratrue sentences express trivial conditions. The step from this to the claim that *no* sentence is ultratrue is relatively small. True, Williamson occasionally uses the methods of normal modal logic in modelling the logic of 'definitely'. Since normal modal logics validate the rule of necessitation—' $\Delta\phi$ ' is a theorem whenever ϕ is—these methods cannot be strictly correct if nothing is ultratrue. But the methods of normal modal logic can be useful tools without being strictly correct, as witness their widespread use in modelling knowledge even by those who reject logical omniscience. Almost none of the philosophical uses to which Williamson puts these methods require taking them any more seriously than this.²⁰

²⁰One exception is Williamson 1999. On the definitions proposed in that paper, for Φ to be "first-order precise", it is not enough for ' $\forall x(\Delta\Phi(x) \vee \Delta\neg\Phi(x))$ ' to be necessarily true: its truth must be 'semantically guaranteed', in a sense that requires it to remain true when prefixed by any number of Δ s. Williamson motivates this by the desire to avoid what he deems 'the counterintuitive situation of higher-order vagueness without first-order vagueness'. But once we realise that no sentence enjoys this kind of "semantic guarantee of truth", we will presumably have to admit, given any sensible definition of '*n*th-order vague', that every sentence is *n*th order vague for some *n*. So if we want anything to count as precise, we had better get used to the situation that Williamson finds counterintuitive.

4 No sentence is ultratrue: second attempt

The foregoing argument from BORDERLINE to MARGIN is highly specific to Williamson's theory. I don't know of any argument for MARGIN that non-epistemicists should find convincing.

Here is a way of thinking about what it would be to accept MARGIN, once we expand it in the obvious way to cover 'dtrue of' as well as 'dtrue'. Where R is some quaternary relation between predicates, communities, worlds and sequences of arguments, let R^* be the quaternary relation that holds, necessarily, among $\Phi, \langle x_1, \dots, x_n \rangle, C$ and w (in that order) iff R holds among the predicate 'dtrue of ... as used by ... at ...' ('dtrue' for short), the sequence of arguments $\langle \Phi, \langle x_1, \dots, x_n \rangle, C, w \rangle$, our community, and the actual world. Say that R is δ -modest iff whenever $R^*(\Phi, \langle x_1, \dots, x_n \rangle, C, w)$, and the distance between w' and w is less than δ , $R(\Phi, \langle x_1, \dots, x_n \rangle, C, w')$. MARGIN entails that the relation of *dtruth* is δ -modest for some positive δ . For MARGIN itself to be *dtrue*, then, it would have to be the case that *every precisification* of 'dtrue' is δ -modest for some positive δ . (Say that an n -ary relation R is a precisification of an n -ary predicate Φ iff 'Necessarily, for all $x_1 \dots x_n$, $R(x_1, \dots, x_n)$ iff $\Phi(x_1, \dots, x_n)$ '⁷ is not dfalse of that relation, as used by us at the actual world.) Once we set epistemicism aside, it is hard to see what aspect of our usage of 'dtrue of' could constrain its precisifications in this way.

By way of contrast, it is easy to see how our usage could impose the weaker constraint that R^* entails R whenever R is a precisification of 'dtrue'. We treat 'definitely' as factive, in the sense that we treat sentences of the form 'if definitely P , then P ' as obvious truths. It is unmysterious how these dispositions could render the sentence 'Whenever "dtrue of" is dtrue of $\langle \Phi, \sigma, C, w \rangle$ as used by us at the actual world, Φ is dtrue of σ as used by C at w ' dtrue in our mouths. But given standard compositional rules, making that sentence come out dtrue requires preventing any R which fails to hold in some cases where R^* holds from being among the precisifications of 'dtrue of'. I don't see any analogous facts about our usage of 'dtrue of' that could, in a parallel way, constrain it to have only δ -modest relations as precisifications.

On the other hand, it is, if anything, even more obvious that if BORDERLINE is true, there is

nothing in our usage of ‘dtrue of’ that could render MARGIN *dfalse* —that could, that is, prevent any of the precisifications of ‘dtrue of’ from being δ -modest for any positive δ . Fortunately, our argument that no sentence is ultratue can be adapted so as to rely on this weaker claim instead of MARGIN. Roughly speaking: since SERIES and MARGIN jointly entail ‘no sentence is ultratue’, the claim that SERIES is dtrue together with the claim that MARGIN is not dfalse jointly entail that ‘no sentence is ultratue’ is not dfalse, which in turn entails that no sentence is ultratue.

Let me restate that argument a bit more carefully, so as to forestall some distracting objections. Let ‘ $M(\delta)$ ’ stand for the claim that δ satisfies the condition specified in MARGIN, that is:

Whenever a sentence S is definitely dtrue as used by us at w , and the distance between w and w' is less than δ , S is dtrue as used by us at w' .

Let ‘ $W(S, n, \delta)$ ’ stand for the claim that there is a sequence $\langle w_0, \dots, w_n \rangle$ which meets the conditions specified in SERIES, that is:

- S1 S is not dtrue as used by us at w_0 .
- S2 w_n is the actual world.
- S3 The distance between w_i and w_{i+1} is always less than δ .
- S4 For any sentence T , necessarily, if $\ulcorner \Delta T \urcorner$ is dtrue as used by us at w_i , then definitely, T is dtrue as used by us at w_i .

The derivation at the beginning of section 3 shows that

(*) If $W(S, n, \delta)$ and $M(\delta)$, then $\ulcorner \Delta^n S \urcorner$ is not dtrue as used by us at the actual world.

is a logical (or at least a mathematical) truth.²¹ As such, (*) is itself dtrue of each $\langle S, n, \delta \rangle$, as used by us at the actual world. But on almost any reasonable account of dtruth for conditionals and conjunctions, a dtrue conditional can have a dfalse consequent only if it has a dfalse antecedent,

²¹While it is controversial whether conditional proof is generally acceptable when vagueness is in question, none of the steps in the argument from section 3 has the features that are supposed to make for failures of conditional proof.

and a conjunction with one dtrue conjunct can be dfalse only if the other conjunct is dfalse. So if (*) and W are both dtrue of $\langle S, n, \delta \rangle$ (as used by us at the actual world) and M is not dfalse of δ , ‘ $\lceil \Delta^n S \rceil$ is not dtrue as used by us at the actual world’ must not be dfalse of $\langle n, S \rangle$. If so, ‘ $\lceil \Delta^n S \rceil$ is dtrue as used by us at the actual world’ is *not* dtrue of $\langle n, S \rangle$. So ‘ $\lceil \Delta^{n+1} S \rceil$ is not dtrue’.

The conclusion that no sentence is ultratrue thus follows from a strengthened version of SERIES together with a weakened version of MARGIN:

SERIES⁺ For each $\delta > 0$ and sentence S , there is an n such that *definitely* $W(S, n, \delta)$.

MARGIN⁻ There is a $\delta > 0$ such that *not definitely not* $M(\delta)$.

The case for SERIES⁺ is not significantly weaker than the case for SERIES: the task of choosing $w_0 \dots w_n$ in such a way that S1–S4 are *definitely* satisfied doesn’t seem especially harder than the task of choosing them in such a way that they are simply satisfied.²² And given BORDERLINE, the case for MARGIN⁻ is quite strong. For MARGIN⁻ to hold, it is sufficient for even *one* of the precisifications of ‘dtrue’ to be δ -modest, for some positive δ . The non-existence of such precisifications would amount to a “penumbral connection” between claims about the dtruth-conditions of arbitrary sentences as used at arbitrary possible worlds and the dtruth-conditions of ‘dtrue’ at the actual world. As such, it would cry out for an explanation in terms of some distinctive feature of our use of ‘dtrue’: in general, when we don’t do anything distinctive to make there be penumbral connections, there aren’t any. One thing we could have done would have been to endow ‘dtrue’ with some precise cutoffs, so that there would be arbitrarily close worlds w and w' such that some sentence

²²The only possible stumbling block is condition S4: it might be thought that the set of worlds of which it is *definite* that ‘definitely’ is used at least as stringently at them as it is at the actual world was too small or scattered to contain a path from actual world to an appropriate w_0 . I doubt that there is anything to this worry: it seems easy to imagine ways of changing the use of ‘definitely’ that would *definitely* either strengthen it or leave it alone. But even if there were, it really wouldn’t matter much, since condition S4 is much stronger than it needs to be for the argument to work. All we really need is that the use of ‘ Δ ’ at each w_i should be similar enough to its actual use for something like MARGIN⁻ to be true: that is, we only need there to be a δ such that for any of the w_i , it is not definitely not the case that when ‘ $\lceil \Delta S \rceil$ is dtrue as used by us at w_i , and the distance between w' and w_i is less than δ , S is dtrue as used by us at w' . It would be straightforward exercise to rewrite the argument using this weaker premise.

S is definitely dtrue as used by us at w and definitely not dtrue as used by us at w' . But according to BORDERLINE, we didn't do this. And we don't seem to have done anything else relevant in the present case.

5 A problem with reference-fixing

This argument for MARGIN⁻ depends essentially on BORDERLINE:

BORDERLINE There is some distance δ such that whenever a sentence S is definitely dtrue as used by some community at w , and the distance between w' and w is less than δ , S is not definitely not dtrue as used by that community at w' .

I motivated BORDERLINE in section 3 by appealing to the idea that the dtruth-conditions of a sentence depend on its use. So long as we focus on such aspects of “use” as people’s dispositions to affirm or deny a sentence under various conditions, and to behave in various ways in reaction to other people affirming and denying it, it will seem obvious that a sufficiently minuscule shift in use could never *definitely* make the difference between a sentence’s being dtrue and its not being dtrue (as evaluated at any given world). However, if we want the claim that dtruth-conditions depend on use to be uncontroversial, we had better make sure to understand “use” more broadly than this—broadly enough so that, for example, the sentence ‘there is water’ counts as being “used in different ways” on Earth and on Twin Earth. And once we pay attention to *this* sort of way for differences in the world to make for differences in the dtruth-conditions of sentences, potential counterexamples to BORDERLINE come quickly to mind.

Suppose we discover that the universe has a finite duration, from Big Bang to Big Crunch. We could then introduce the expression ‘aeon’ by issuing a stipulation: ‘Let “one aeon” name the duration of the universe.’ We thereby, let’s suppose, introduce a precise expression which is definitely, necessarily, dtrue of all and only those temporal intervals whose duration is the same as the actual duration of the universe. Thus the sentence ‘The universe lasts for at most one aeon’ is definitely dtrue as used by us at the actual world. But this same sentence is definitely not dtrue

(at the actual world) as used by us at a close world where the Big Crunch happens a little earlier. Since we may suppose that the laws of nature allow such worlds to be arbitrarily close to the actual world, this is a counterexample to BORDERLINE.²³

There are ways of fixing up our argument to make it proof against these counterexamples. The most obvious strategy is to put some *ad hoc* restriction into BORDERLINE—something like ‘... so long as the use of S at w and w' doesn't involve a reference-fixing description which definitely denotes different things w and w' ’—and to use this to argue for a correspondingly restricted version of MARGIN⁻. We could still argue for the conclusion that no sentence whatsoever is ultratrue, by appealing to the claim that there are some sentences—‘ $0 \neq 1$ ’, say—which are ultratrue if anything is, and whose use doesn't involve the kind of reference fixing that makes for exceptions to BORDERLINE.²⁴

Still, it is interesting to see whether we can find any non-*ad hoc*, defensible principles in the vicinity of BORDERLINE and MARGIN⁻. In the remainder of this section I will discuss two possible strategies for formulating such principles.

The first strategy is to understand ‘ S is dtrue as used by C at w ’ in a way modelled on what the tradition of two-dimensional semantics calls the “primary intension” as opposed to the “sec-

²³Real-world uses of reference-fixing descriptions don't pose any obvious problems for BORDERLINE. ‘One kilogram’ is stipulated to be the rest mass of a particular platinum-iridium cylinder: but because of fluctuations over time in the mass of the cylinder, vagueness as regards the locations of its boundaries, and perhaps also further quantum-field-theoretic sources of vagueness in claims about the masses of particular material bodies, it is not plausible that this stipulation has made ‘one kilogram’ perfectly precise, and thus not plausible that there is any possible object of which ‘has a mass of at least a kilogram’ is definitely dtrue as used by us at the actual world and definitely not dtrue as used by us at worlds arbitrarily similar to the actual world. ‘One second’ may for all I know be completely precise, but unlike our envisaged definition of ‘aeon’, its definition (p. 10 above) doesn't seem to depend on anything nomologically contingent, so the ‘dtruth’-conditions of ‘one second’ wouldn't vary between the actual world and other nomologically possible worlds where it is associated with the same description. Expressions constructed using ‘actually’ are the best real-world candidates to be counterexamples to BORDERLINE.

²⁴In fact, we need something a bit stronger than this: we need an assurance that such reference-fixing is not a feature of the usage of any of the sentences ‘ $\Delta^i S$ ’, and we need this to be true not only at the actual world but at all of the worlds in some appropriate sequence $\langle w_0, \dots, w_n \rangle$. But this seems fine: none of the sentences ‘ $\Delta^i 0 \neq 1$ ’ seems to be anything like ‘The universe lasts for at least one aeon’ in the relevant respects.

ondary intension” of S . The idea is that while differences in the denotation of a reference-fixing description make a difference to the secondary intension of an expression, they make no difference to its primary intension, which is always the same as the primary intension of the reference-fixing description. Thus, for any world w where the expressions ‘one aeon’ and ‘the duration of the universe’ are associated in the right way, a necessary truth that ‘the universe lasts for at most one aeon’ is dtrue in the primary sense as used by us at w if the duration of the universe is finite.

The *two-dimensionalist* programme (Chalmers & Jackson 2001; Chalmers 2006) attempts to assign distinct primary and secondary intensions to many ordinary expressions, in such a way that the primary intension of each expression encodes certain facts about its epistemological properties. The application of two-dimensional machinery we are presently contemplating doesn’t require anything so ambitious. The main objection to assigning an expression like ‘water’ a primary intension distinct from its secondary intension is the sheer difficulty of finding any non-arbitrary principle for reading a primary intension off the use of this expression, let alone one that captures anything epistemologically significant. But ‘aeon’ poses a problem for BORDERLINE precisely because the description ‘the duration of the universe’ plays such a clear and non-arbitrary role in regulating its use. If its use suffered from the sort of messiness that makes the project of assigning interesting primary intensions to most expressions so hard, it would no longer be plausible that an arbitrarily small change could make the difference between definite dtruth and definite lack of dtruth (at the actual world, in our old, “secondary” sense). Thus, even a very modest dose of two-dimensionalism, on which primary and secondary dtruth-conditions diverge only in the rare cases where it is completely clear how to draw the distinction in a non-arbitrary way, should be enough to make BORDERLINE proof against the counterexamples we have been considering in this section.

This first strategy will, however, lead us into new difficulties if we combine it with the meta-linguistic analysis of $\ulcorner \Delta S \urcorner$ as $\ulcorner \text{‘}S\text{’ is dtrue as used by us at the actual world} \urcorner$. Given the standard treatment of the primary intensions of sentences involving ‘actual’ and ‘actually’, this analysis makes the primary dtruth-conditions of $\ulcorner \Delta S \urcorner$ come apart from those of S in a surprising way. For example, consider a world w where we use ‘cat’ the way we actually use ‘lemur’, while using all

other words just as we actually do. Since there are fewer than a million lemurs, ‘There are more than a million cats’ is *dfalse* (in both the primary and secondary senses) as used by us at w . But if ‘Definitely, there are more than a million cats’ is synonymous at w with ‘“There are more than a million cats” is *dtrue* as used by us at the actual world’, it is *dtrue* in the primary sense, since the use of ‘“There are more than a million cats” is *dtrue* as used by us at the actual world’ doesn’t vary between w and the actual world in any relevant way. Unsurprisingly, a notion of *dtruth* which made it this easy for ‘ ΔS ’ to be *dtrue* without S being *dtrue* would make a mess of our argument.²⁵

At this point, the proponent of the ‘actually’-involving analysis of ‘definitely’ might consider simply eliminating all occurrences of ‘definitely’ from the argument in favour of the putatively more fundamental ‘*dtrue* as used by ... at ...’. There may be a workable argument to be found here, but things quickly get complicated, as we have to deal with claims about the *dtruth* of predicates of sequences of arguments inside which other sequences may be deeply nested. The task of formulating a compelling principle strong enough to play the role of *BORDERLINE* in such an argument is quite challenging.

The second strategy is to rethink our stipulation that claims of the form ‘ S is *dtrue* as used by C at w ’ are to be evaluated with respect to the actual world. As I pointed out (comment (a) in section 3) the argument would be hopeless if we had instead understood ‘*dtrue* as used by C at w ’ as equivalent to ‘*dtrue* at w as used by C at w ’: in that case, *BORDERLINE* (and *MARGIN*, and *MARGIN*⁻) would have been immediately refuted by the existence of precise nomologically contingent sentences like ‘The duration of the universe is at least n seconds.’ But we can state a principle in the spirit of *BORDERLINE* that allows for such expressions, by trading in our simple measure of

²⁵The problem turns out to lie with S4, which becomes extremely demanding if ‘ Δ ’ is analysed using ‘actually’, ‘actually’ is given the usual two-dimensional semantics, and ‘*dtrue*’ is understood in the primary sense. Suppose that ‘ Δ ’ and ‘*dtrue* as used by us at the actual world’ are used in the same way at w_i as they are at the actual world. Then whenever ‘ ΔT ’ is *dtrue* in the primary sense as used by us at the actual world, it is *dtrue* in the primary sense as used by us at w_i . If w_i also satisfies S4, it follows that whenever ‘ ΔT ’ is *dtrue* as used by us at the actual world, it is definite that T is *dtrue* in the primary sense as used by us at w_i . For this to be true, our use of T at w_i would have to be quite similar to our use of T at the actual world. Given this, there is no longer any clear reason to expect to be able to get from the actual world to an appropriate w_0 via a chain of worlds all of which satisfy S4.

similarity between worlds for a more complicated measure of similarity between ordered triples of worlds, communities and sentences:

BORDERLINE* There is a $\delta > 0$ such that whenever a sentence S is definitely dtrue at w as used by C at w , and the distance between $\langle S, C, w' \rangle$ and $\langle S, C, w \rangle$ is less than δ , S is not definitely not dtrue at w' as used by C at w' .

How would the distance metric have to work for **BORDERLINE*** to be defensible? Clearly, if S is ‘The universe lasts for at least n seconds’, and C is a community that uses ‘second’ much as we do at both w and w' , the fact that the universe lasts for less than n seconds at w and more than n seconds at w' must be sufficient for the distance between $\langle S, C, w \rangle$ and $\langle S, C, w' \rangle$ to exceed some positive threshold, no matter how close w and w' might be in our old sense. We can achieve this, I think, by thinking of the similarity relation between ordered triples as grounded, at least in part, in more or less *natural* relations, in the same way that the similarity relations among objects are generally thought of as grounded in their natural properties (see Lewis 1983). The more natural a relation, the more the fact that it holds between S , C and w but not between S , C and w' will contribute to dissimilarity between $\langle S, C, w \rangle$ and $\langle S, C, w' \rangle$. As a special case, the more natural a function f from sentences, communities and worlds to propositions, the more the fact that $f(S, C, w)$ holds at w while $f(S, C, w')$ fails to hold at w' will make for dissimilarity. If C employs the actual definition of ‘second’ involving the number 9,192,631,770 at both w and w' , then the fact that S is true at w and false at w' on an interpretation on which ‘one second’ stands for 9,192,631,770 units of some fairly natural duration will make for substantial dissimilarity between $\langle S, C, w \rangle$ and $\langle S, C, w' \rangle$. If, instead, making S have different truth values at w and w' required interpreting ‘one second’ as standing for 9,162,631,771 units of the same duration, this would make for much less dissimilarity between the triples.

If you can make sense of the background ideology of degrees of naturalness, you should find **BORDERLINE*** plausible. It is an instance of a plausible general schema, which captures the idea that we can only achieve precision along some dimension when the dimension contains sufficiently natural joints for our use to pick up on:

(N) There is a $\delta > 0$ such that whenever it is definitely the case that $F(x_1, \dots, x_n)$, and the distance between $\langle x_1, \dots, x_n \rangle$ and $\langle y_1, \dots, y_n \rangle$ is less than δ , it is not definitely not the case that $F(y_1, \dots, y_n)$.

The picture is that by working hard—e.g. by formulating long and complicated definitions—we can make the use of a predicate sensitive to less and less natural distinctions, thereby reducing the maximum value of δ for the predicate in question. But we can only do a finite amount of this kind of work. If it turns out that objects of which some predicate is dtrue can be arbitrarily close to objects of which it is dfalse on some similarity measure, the measure in question must fail to represent all the natural joints in the relevant space.

Can we use BORDERLINE* to argue that no sentence is ultratrue? The argument from BORDERLINE to MARGIN⁻ in section 4 can be adapted to yield an analogous argument for

MARGIN* There is a $\delta > 0$ such that it is not definitely not the case that: whenever a sentence S is definitely dtrue at w as used by C at w , and the distance between $\langle S, C, w \rangle$ and $\langle S, C, w' \rangle$ is less than δ , S is dtrue at w' as used by C at w' .

But how are we to get from MARGIN* to the conclusion that no sentence is ultratrue? The space of $\langle S, C, w \rangle$ triples will turn out to be far from continuous under any metric that tracks natural properties and relations. On any reasonable way of thinking about degrees of naturalness, it is inevitable, leaving aside perfectly symmetric universes, that for *any* distinct triples $\langle S, C, w \rangle$ and $\langle S', C', w' \rangle$, there will be a relation R that has *some* positive degree of naturalness such that $R(S, C, w)$ but not $R(S', C', w')$. Thus, if we want to trace a path from $\langle S, C, w \rangle$ to $\langle S', C', w' \rangle$ in small steps, there will always be some lower bound to the size of steps we can allow ourselves. A space of possibilities with a metric based on naturalness is like a fractal landscape, crosscut so thoroughly with cracks that one can never get anywhere without stepping over a crack of some nonzero width.

This makes it hard to state a premise that can take over the role of SERIES or SERIES⁺ in the new framework. But I don't think that this problem is too serious. Even without anything like an articulated metasemantics, we can see that the power of naturalness to make for definite cutoffs in

the extensions of our predicates falls off quite quickly. The property *having a mass of more than 45955882 Planck masses* is *somewhat* natural, but it is not natural enough for it to be at all plausible that ‘has a mass of at least kilogram’ as used in, say, 1800 was definitely dtrue of all and only the things whose mass was more than 4595582 Planck masses. This suggests that some δ_0 meets the condition specified in MARGIN* and is fairly big—big enough for the set of all $\langle S', C', w' \rangle$ reachable from any given $\langle S, C, w \rangle$ by way of steps no bigger than δ_0 to be, in general, quite extensive. So the following principle has some plausibility:

SERIES* For each sentence S , there is a positive δ_0 meeting the condition specified in MARGIN*, such that for some n , definitely, there is a sequence $\langle w_0, \dots, w_n \rangle$ for which:

- S1 S is not dtrue at w_0 as used by us at w_0 .
- S2 w_n is the actual world.
- S3 For each i , the distance between $\langle \ulcorner \Delta^i S \urcorner, \text{us}, w_i \rangle$ and $\langle \ulcorner \Delta^i S \urcorner, \text{us}, w_{i+1} \rangle$ is less than δ_0 .
- S4 For any sentence T , necessarily, if $\ulcorner \Delta T \urcorner$ is dtrue as used by us at w_i , then definitely, T is dtrue as used by us at w_i .

Given SERIES* and MARGIN*, we can argue that no sentence is ultratrue in the same way as before.

If you thought that the relation S *being dtrue at w as used by C at w* was itself fairly natural, you would have no reason to accept SERIES* (unless you already believed for some other reason that nothing was ultratrue). For in that case, if we started out assuming that S was ultratrue as used by us at the actual world, the stipulation that S is not dtrue as used by us at w_0 would be enough to entail that for some i , the triples $\langle \ulcorner \Delta^i S \urcorner, \text{us}, w_i \rangle$ and $\langle \ulcorner \Delta^i S \urcorner, \text{us}, w_{i+1} \rangle$ differ as regards the fairly natural relation *being dtrue as used by...at...*, and thus count as fairly far apart on the relevant similarity measure.²⁶ But let’s assume the more orthodox view that takes physics to be

²⁶Hawthorne’s (2006) suggestion that ‘true’ (as opposed to ‘dtrue’) expresses a natural relation would also, I think, undermine the plausibility of SERIES*; but the relevant considerations in that case are more intricate.

our best guide to the structure of natural properties. In that case, SERIES* seems quite secure for many values of S —‘ $0 \neq 1$ ’, for example. It is not plausible, from the physicalistic perspective, that any of the remotely natural relations that hold between any sentence ‘ $\Delta^i 0 \neq 1$ ’, us, and the actual world is even *sufficient* for the sentence to be dtrue as used by us at the actual world. The physical facts about us that make some of these sentences dtrue in our mouths are just too complicated. So it should be possible to find a sequence $\langle w_0, \dots, w_n \rangle$ satisfying conditions S1, S2 and S4, and such that whenever an even remotely natural relation holds between ‘ $\Delta^i 0 \neq 1$ ’, us, and w_i , it also holds between ‘ $\Delta^i 0 \neq 1$ ’, us, and w_{i+1} . If so, we can make the distance between $\langle \Delta^i S, \text{us}, w_i \rangle$ and $\langle \Delta^i S, \text{us}, w_{i+1} \rangle$ very small, on a naturalness-respecting metric. We will still have to step across tiny cracks, corresponding to relations which aren’t even “remotely” natural. But an account of dtruth that allowed even these very minor joints in nature to endow vague predicates with definite cutoffs would, I think, grossly overestimate the role of naturalness in metasemantics.²⁷

²⁷Wait: isn’t *being actualised* (understood in such a way that it is contingent which world is actualised) a highly natural property? And isn’t *being an S , C and w such that w is actualised* therefore a highly natural relation? And won’t it follow, therefore, that whenever w is the actual world and w' isn’t, the distance between $\langle S, C, w \rangle$ and $\langle S, C, w' \rangle$ is fairly large, so that there can be no sequence of the kind required by SERIES*?

The point is well taken: the discussion in the main text implicitly assumes that the relevant distance relation is itself a necessary one, for which only necessary natural relations need to be taken into account. One way of fixing up the argument is to change BORDERLINE* to make it explicit that only necessary relations are to be taken into account. So strengthened, BORDERLINE* will no longer be an instance of a plausible general schema: for (N) to be plausible for ordinary contingent predicates like ‘is positively charged’, the distance relation must of course take contingent natural relations into account as well as necessary ones. But the stronger version of (N) does seem to be plausible for many predicates whose extensions are a necessary matter, for example predicates of the form ‘ x is F at w ’. The exceptions are predicates that achieve precision using devices like ‘actually’, like ‘ x is the same length as it actually is at w ’, and similar predicates introduced using reference fixing. But it is hard to see any relevant similarity between ‘dtrue’ and these.

Another response is simply to drop the requirement that w_n be the actual world, and replace it with the requirement that if $\Delta^n S$ is dtrue at the actual world, it is dtrue at w_n . This seems safe—surely, if there are ultratrue sentences, there are sentences which would still have been ultratrue if things had been very slightly different. And once we no longer have to take a step from a non-actual world to the actual world, there is no obvious further reason why taking contingent natural relations into account in the distance metric should undermine the case for SERIES*.

Not all sentences are like ‘ $0 \neq 1$ ’ in this respect. If the number of particles is finite, then there will be sentences which specify the position and momentum of each particle, using physically natural units, up to any desired degree of accuracy. If so, it will sometimes happen that a (ridiculously long!) sentence S has a highly natural interpretation on which it entails that it itself is dtrue as used by us. If so, it is also true of each sentence $\Delta^i S$ that it has a highly natural interpretation on which it entails that S is dtrue as used by us: we can always just interpret Δ as a vacuous operator. In that case, if S is in fact dtrue as used by us, there will be no way to choose $\langle w_0, \dots, w_n \rangle$ so as to satisfy S3: for some i , the triples $\langle \ulcorner \Delta^i S \urcorner, \text{us}, w_i \rangle$ and $\langle \ulcorner \Delta^i S \urcorner, \text{us}, w_{i+1} \rangle$ will have to differ as regards whether their first member is true at their third member on the highly natural interpretation in question.

So SERIES* cannot be defended in full generality. But we can still argue that no sentence is ultratrue indirectly, by first arguing, say, that ‘ $0 \neq 1$ ’ is not ultratrue, and then arguing that if any sentence were ultratrue, ‘ $0 \neq 1$ ’ would have to be. Or we could argue, first, that no sentence below some given length is ultratrue, and second, that since every sentence has short sentences among its logical consequences, if any sentence were ultratrue, some short sentence would have to be. This is a bit disappointing: it would be nicer to be able to argue by appealing to completely general premises. Still, for those who can stomach the hefty dose of vagueness involved in all this talk about degrees of naturalness and similarity relations that respect them, the present approach has dialectical advantages, in that it lets us avoid the task of formulating and defending an *ad hoc* restriction of MARGIN⁻ that avoids counterexamples involving reference-fixing.

6 How far does definiteness iterate?

Once we have agreed that there is a finite i such that no sentence beginning with Δ^i is dtrue, it is natural to ask what the smallest such i is.²⁸ This much is clear: the vaguer we take ‘definitely’ to be, the fewer iterations we should expect it to sustain. But what sorts of numbers are we talking about?

²⁸Opponents of classical logic may reject the presupposition that there is a smallest such i . Still, they will want to know how to answer various specific questions of the form ‘Are there dtrue sentences starting with Δ^i ?’

I don't know how to make progress with this question in a way that is neutral between different theories of vagueness. So I will approach it from the standpoint of the following simple metalinguistic theory (Lewis 1969; Dorr 2003): for S to be dtrue as used by C at w is for there to be some true proposition P , such that there prevails among C at w a system of conventions that permits asserting S while believing P , and forbids asserting the negation of S while believing P . For short, let's say in this case that S is 'conventionally favoured' by C at w .

Those who have doubts about the very notion of a linguistic convention may worry that no sentence whatsoever—not even

C0 $0 \neq 1$

—will get to be dtrue on this analysis. But if we bracket these doubts, I see no special reason to doubt that among the sentences that are conventionally favoured as used by us (readers of this paper) at the actual world, there will be some which characterise other sentences as “conventionally favoured”—for example,

C1 C0 is conventionally favoured by us at the actual world.

If our community has any linguistic conventions at all, then someone who insisted on asserting the negation of C1 while believing all relevant truths about how we treat C0 would surely manifest a failure to abide by them (including the conventions concerning the uses of the expressions 'C0' and 'conventionally favoured' that I have just instituted). Just as Yul Brynner is a paradigm case of the sort of person to whom it would be appropriate to apply 'bald', so $\langle C0, \text{us, the actual world} \rangle$ is a paradigm case of the sort of triplet to which it would be appropriate to apply 'conventionally favoured by...at...'. And if dtruth is conventional favouredness, the claims of sentences like 'C0 is dtrue as used by us at the actual world' and ' $0 \neq 1$ ' to count as conventionally favoured, and thus dtrue, seem as strong as those of C1.

The case that

C2 C1 is conventionally favoured by us at the actual world

is conventionally favoured by us at the actual world is not quite as strong, but is still compelling. Granted, no-one would be tempted to use $\langle C1, \text{us, the actual world} \rangle$ as a paradigm case in introducing someone to the use of ‘conventionally favoured’: by comparison to our use of mathematical vocabulary, our use of words like ‘convention’ is fluid, even anarchic. You can imagine someone being so impressed by this contrast that they insisted on applying the term ‘convention’ only to those regularities in linguistic activity that attained the level of rigidity encountered in domains such as mathematics, and thus was disposed to assert the negation of C2. But this would be an excessively finicky way to use ‘convention’. It would surely not be in accord with our actual conventions concerning the use of that word, at least in the kinds of contexts we are presently concerned with.

This case becomes much harder to make when we turn to

C3 C2 is conventionally favoured by us at the actual world.

The level of finickiness in applying the word ‘convention’ that might lead someone to deny C3 is not nearly so high. Without actually being disposed to deny C2, one could still be exacting enough to insist that someone who *did* act on such a disposition would not thereby count as “violating a convention” about the use of the word ‘convention’. This lower level of finickiness is considerably easier to feel sympathy with; it is less alien to our ordinary practice in applying the word ‘convention’. Would even this constitute a failure to abide by the conventions concerning the use of ‘convention’? At this point I don’t feel at all sure what to say. And I feel even less sure that if I were to answer ‘no’, I would thereby be violating any convention.

Given the account of dtruth as conventional favouredness, doubts about whether C3 is conventionally favoured will carry over to sentences like ‘ $\Delta^3 0 \neq 1$ ’. A disposition to deny ‘ $\Delta^3 0 \neq 1$ ’ could arise from the combination of a degree of finickiness about the use of ‘convention’ sufficient to prompt the denial of C3 with an explicit endorsement of the theory that to be dtrue is to be conventionally favoured. It would thus be hard for a proponent of that theory to claim that ‘ $\Delta^3 0 \neq 1$ ’ is conventionally favoured while denying that C3 is.

So if we adopt a metalinguistic theory of vagueness along these lines, we will find it hard to

maintain that $\lceil \Delta^3 S \rceil$ is dtrue for any S . This doesn't mean that we will find it easy to argue that it *isn't* ever dtrue: the point at which there starts to be overall theoretical pressure to claim that $\lceil \Delta^i S \rceil$ is *not* dtrue comes a bit later.²⁹ But in view of the sharp drop between the degrees of finickiness required to prompt the assertion of $\lceil \neg \Delta^i S \rceil$ and $\lceil \neg \Delta^{i+1} S \rceil$ that emerged in the cases we examined, I can't see how it could come *much* later. I think a fairly compelling argument could be made from the theory of dtruth as conventional favouredness to the claim that ' $\Delta^5 0 \neq 1$ ' is not dtrue.

I am not sure to what extent these considerations carry over to theories different in character from the convention-based theory. But one general point can be made. If we estimate the "degree of vagueness" of expressions like 'borderline', and 'definitely' by comparing their use to the use of other expressions in the language, the natural conclusion is that they are *extremely* vague. The use of these expressions outside of philosophy is largely restricted to a few rather stylised contexts. Whereas our mastery of words like 'causes' or 'believes' involves impressive feats of co-ordination whose inner workings are far from being transparent to us, our ability to use expressions like 'borderline' is not that big an achievement. There is conspicuously little discipline of the kind that is generally required for the range of cases where a vague expression is dtrue to outstrip the range of cases where its application is uncontroversial. Given these facts, merely admitting that a sentence like ' $\Delta^5 0 \neq 1$ ' is controversial is already enough to put some pressure on the claim that it is dtrue.

7 Why does it matter?

The question whether any sentences are ultratrue is not of merely technical interest. Its answer bears on several issues at the heart of the philosophy of vagueness.

To begin with: if no sentence is ultratrue, there is just no sense in which the rule of " Δ -introduction", $P \vdash \Delta P$, is valid. The claim that Δ -introduction is valid has had a wide following in the literature on vagueness (e.g. Fine 1975; Wright 1987; Heck 1993; Keefe 2000; Field 2000,

²⁹And of course the point at which there starts to be theoretical pressure to claim that $\lceil \Delta^i S \rceil$ is *dfalse* comes later still.

2003a). It provides one important motivation for thinking that vagueness requires revision of classical logic. For the notion of definiteness would be pointless if $\neg\Delta P \vdash \neg P$ were valid; but given the classical metarule of proof by contradiction, the latter rule must be valid if Δ -introduction is. It also constitutes a serious obstacle to the project of analysing borderlineness and definiteness. If Δ -introduction is valid, any putative analysis of ‘definitely’ faces something like the open question argument: ΔP follows from P ; the putative analysis of ΔP does not; so the putative analysis must be incorrect.

The conclusion that no sentence is ultratrue undercuts any temptation to count Δ -introduction as valid. Assuming that validity in the relevant sense is transitive, all inferences of the form ‘ $P \vdash \Delta^n P$ ’ must be valid if Δ -introduction is. But even without proof by contradiction, the idea that something dfalse can be validly derived from every sentence cannot be taken seriously. The production of a valid argument from a sentence we have uttered assertively to something we regard as dfalse should, at the very least, compel us to retract or qualify our assertion. But it would be absurd to react to the realisation that no sentence is ultratrue by abandoning the practice of making assertions.

For the purposes of arguing against Δ -introduction, the claim that very few sentences are ultratrue would do as well as more sweeping claim that none are, since the claim that *almost* all sentences have dfalse logical consequences is not significantly less absurd than the claim that all do. But the more sweeping claim still has important ramifications.

While nobody would be tempted to think of the identification of the ultratrue sentences *per se* as a central goal of philosophy, the idea that the set of ultratrue sentences is small but nonempty is defensible only on the assumption that there is some more philosophically important feature which the ultratrue sentences share, and which explains why just they get to be ultratrue. This feature might be analyticity (though this would require an unusually restrictive conception of analyticity); it might be logical truth, conceived of as something less arbitrary than mere truth (or analytic truth, or necessary truth) in virtue of the meaning of some list of “logical constants”. Whatever the deep property that explains ultratruth is, philosophers will naturally be drawn to it as a standard of ultimate theoretical success. We will aspire not merely to express interesting truths (perhaps necessary

truths) about the subject matters we investigate (causation, knowledge, right and wrong...): we will aspire to uncover the “logics” of these domains, in some exalted sense. And to the extent that we adopt such a goal, we will see a difference in kind between our inquiry and inquiry in domains where ultratruth is not on the cards. But if no sentence is ultratrue, this way of distinguishing philosophical inquiry from inquiry of other kinds can be dismissed as chimerical. There may still be sharp categorical distinctions between different kinds of *facts* (or propositions); and philosophy may differ from other fields in the kinds of facts it aims to identify. But if there are no sharp discontinuities in the space of ways in which sentences can be used by communities, there can be no categorical distinction between the *sentences* philosophers aim to produce and sentences of other kinds. When it comes to putting our thoughts into words, we must all muddle along in the same way, doing our best to make ourselves understood with the limited verbal tools at our disposal.

References

- Barbour, J 2006, *The End of Time*, Oxford University Press, Oxford.
- BIPM 2006, *SI Brochure*, Bureau Internationale des Poids et Mesures, 8th edn., URL http://www.bipm.org/en/si/si_brochure.
- Braun, D & Sider, T 2007, 'Vague, so untrue', *Noûs*, vol. 41, pp. 133–156.
- Chalmers, D 2006, 'Two-dimensional semantics', in E Lepore & BC Smith (eds.) *The Oxford Handbook of Philosophy of Language*, Oxford University Press, Oxford.
- Chalmers, D & Jackson, F 2001, 'Conceptual analysis and reductive explanation', *Philosophical Review*, vol. 110, pp. 315–361.
- Dorr, C 2003, 'Vagueness without ignorance', in J Hawthorne & D Zimmerman (eds.) *Philosophical Perspectives 17: Language and Philosophical Linguistics*, pp. 83–113, Blackwell, Oxford.
- Field, H 2000, 'Indeterminacy, degree of belief, and excluded middle', *Noûs*, vol. 34, pp. 1–30.
- Field, H 2003a, 'No fact of the matter', *Australasian Journal of Philosophy*, vol. 81, pp. 457–480.
- Field, H 2003b, 'The semantic paradoxes and the paradoxes of vagueness', in J Beall & M Glanzberg (eds.) *Liars and Heaps*, pp. 262–311, Oxford University Press, Oxford.
- Field, H 2008, *Saving Truth from Paradox*, Oxford University Press, Oxford.
- Fine, K 1975, 'Vagueness, truth and logic', *Synthese*, vol. 30, pp. 265–300.
- Graff, D 2002, 'An anti-epistemicist consequence of margin for error semantics for knowledge', *Philosophy and Phenomenological Research*, vol. 64, pp. 127–142.
- Hawthorne, J 2006, 'Epistemicism and semantic plasticity', in *Metaphysical Essays*, Oxford University Press, Oxford.

- Heck, RG, Jr. 1993, 'A note on the logic of (higher-order) vagueness', *Analysis*, vol. 53, pp. 201–208.
- Keefe, R 2000, *Theories of Vagueness*, Cambridge University Press, Cambridge.
- Lewis, D 1969, *Convention: A Philosophical Study*, Blackwell, Oxford.
- Lewis, D 1983, 'New work for a theory of universals', *Australasian Journal of Philosophy*, vol. 61, pp. 343–377.
- Prior, A 1968, 'Changes in events and changes in things', in *Papers on Time and Tense*, pp. 1–14, Oxford University Press, Oxford.
- Quine, Wv 1940, *Mathematical Logic*, Harvard University Press, Boston.
- Soames, S 1998, 'The modal argument: Wide scope and rigidified descriptions', *Noûs*, vol. 32, pp. 1–22.
- Sorensen, RA 1985, 'An argument for the vagueness of "vague"', *Analysis*, vol. 45, pp. 134–137.
- Williamson, T 1987, 'On the paradox of knowability', *Mind*, vol. 96, pp. 256–261.
- Williamson, T 1994, *Vagueness*, Routledge, London.
- Williamson, T 1999, 'On the stucture of higher-order vagueness', *Mind*, vol. 108, pp. 127–143.
- Williamson, T 2000, *Knowledge and its Limits*, Oxford University Press, Oxford.
- Wright, C 1987, 'Further reflections on the sorites paradox', *Philosophical Topics*, vol. 15, pp. 227–290.