

Vagueness as Indeterminacy

Brian Weatherson, Rutgers University/Arché, University of St Andrews*

July 20, 2008

Recently there has been a flurry of proposals on how to ‘define’ vagueness. These proposals are not meant to amount to *theories* of vagueness as, for instance, epistemic or supervenient theories of vagueness are. That is, they are not meant to provide solutions to the raft of puzzles and paradoxes traditionally associated with vagueness. Rather, they are meant to give us a sense of which terms in the language are vague, and to use Matti Eklund’s phrase, in what their vagueness consists. Doing this *might* be a prelude to a successful theory of vagueness, or it might just be an interesting classificatory question in its own right.

When this activity started, most notably with Patrick Greenough’s “A Minimal Theory of Vagueness”, I suspected that it would be a hopeless project. Imagine, I thought, trying to give a definition of what causation is that didn’t amount to a theory of causation. That project seems hopeless, and I didn’t think the prospects for a definition of vagueness were much better. I now think I was wrong, and we can learn a lot from thinking about which terms are vague, independent of our theory of vagueness. (As we’ll get to below, Greenough’s theory isn’t marketed as a definition of vagueness, but rather a ‘minimal theory’ to which all parties can agree. But it has been taken, e.g. by Eklund and Nicholas Smith, to be providing a rival to genuine definitions of vagueness, and I’ll follow Eklund and Smith in this respect.)

The point of this exercise is not to give an analysis of how the man on the Clapham omnibus uses ‘vague’ and its cognates. As is widely recognised, ‘vague’ is often used in ordinary language as a predicate that applies to claims like *The Grand Canyon is between 2 and 2 trillion years old*, i.e. claims that are consistent with a wide range of possible worlds. That’s not the sense of ‘vague’ which philosophers use, nor the sense we are trying to define. But nor should we think we are just trying to analyse philosophical use of ‘vague’. The philosophers’ usage may be our starting point, but if we find philosophers have traditionally been ignoring theoretically important commonalities, or blurring theoretically important distinctions, our best definition may well amount to a revision of philosophical usage.

The game, I think, is one of setting goals for what a theory of vagueness should do. It is a legitimate objection to a theory of vagueness that it isn’t comprehensive, that it doesn’t cover the field. If supervenientism was only a theory of how vague words that started with consonants behaved, for example, that would be a problem for supervenientism. But to press objections of this form we must have an antecedent answer to the question of which words are in the field, and hence should be covered. That’s the good question which these definitions of vagueness address. Because I take this to be the important issue, I’m going to start this paper with a bunch of

* Thanks to David Chalmers, Matti Eklund, Delia Graff Fara, Patrick Greenough, Mark Johnston, Daniel Nolan, Nicholas J. J. Smith, Andrew Strauss and audiences at ANU, Princeton and Arché.

examples of apparently vague, and apparently non-vague, terms. We'll then look at which theories do the best job at systematising intuitions about these cases. I'll then argue that the best way to systematise our intuitions about these cases while respecting theoretically important commonalities and distinctions is to take vagueness to be indeterminacy, while staying silent for now on whether the indeterminacy is semantic or epistemic. In doing so I'm returning to a traditional view of vagueness, one that is discussed in such classic works as Kit Fine's statement of supervaluationism (Fine 1975). So I make no claim to originality in my conclusions here, though I hope at least *some* of the arguments are original.

1. Examples

I'm going to introduce five classes of examples, which will serve as our data in what follows. I'll give a fairly tendentious description of each class to orient us before starting. Our five classes are (a) words that are indeterminate but not vague, (b) vague words that are not predicates, especially predicate modifiers, (c) vague predicates whose conditions of application are contentious, (d) vague predicates whose application depends on discrete states of the world, and (e) vague predicates that do not determine boundaries.

1a. Indeterminacy without Vagueness

Many philosophers, if asked, would say that vague words are those that have borderline cases. As noted above, Fine (1975) takes exactly this view. My preferred view, that vagueness is indeterminacy, is a simple generalisation of this view to non-predicates. But it is a commonplace of the literature on definitions of vagueness that this won't do because of examples of indeterminacy without vagueness. Two examples are commonly used. One of these is Sainsbury's example *child** (Sainsbury 1991). By definition, the extension of *child** is the set of persons under sixteen years old, and its anti-extension is the set of persons eighteen years old or older. Sixteen and seventeen year olds are borderline cases. The intuition is that even though *child** has borderline cases it is not vague, because there are sharp boundaries to its borderline.

A similar case arrives with *mass* as it is used by a Newtonian physicist. (I'm grateful to Delia Graff Fara for pointing out the connection here.) As Field (1973) showed, *mass* is indeterminate between two meanings, rest mass and proper mass. But it is intuitively *not* vague, because it is determinate that it means either rest mass or proper mass. These cases are well discussed in the existing literature, and I won't say much more about them here, save to note that one of the examples that is usually taken to be very problematic for the vagueness as indeterminacy view, *child**, is not synonymous with any term in any natural language. This is not a reason that it could not serve as a counterexample, because a definition should cover all terms actual and possible.

1b. Non-Predicate Vagueness

Not only predicates are vague. There is an extensive literature on vague singular terms. Arguably many determiners are vague. And, as I'll stress here, many predicate modifiers are vague.

We can make an intuitive distinction between vague and precise predicate modifiers. Compare the following two (obviously artificial) predicate modifiers. (I owe these examples to David Chalmers.) Where F is a predicate such that Fa is true iff for some variable v , $v(a) > x$ and v has a natural zero value (e.g. like height and unlike utility) then we can define *doubly F* and *bigly F*. It is true that a is doubly F iff $v(a) > 2x$ and a is bigly F iff $v(a)/x$ is big. Now there's a good sense in which *doubly* is a precise modifier, for the modification it makes to its attached predicate can be precisely defined, while *bigly* is a vague modifier. That's the sense in which I mean some modifiers are vague and others are precise. Note that even though *doubly* is precise it can be a constituent of a vague predicate, such as *doubly tall*. That makes sense; just as a vague sentence need only contain one vague word, so need a vague complex predicate need only contain one vague word.

Now we might well ask whether natural language modifiers like *very* are vague or precise. I'm sad to say that I really don't have an answer to that question, but I think it's an excellent question. To get a sense of how hard it is, note one awkward feature of *very* – it is most comfortable attaching to words that are themselves vague. For instance (1a) is a sentence of English while (1b) is not.

- (1) a. Jack is very old.
 b. *Jack is very forty-seven years old.

I don't know whether this is a universal feature of *very*. My best guess is that it is though in conversation some people have proposed interesting putative counterexamples. (I'm grateful here to Daniel Nolan.) But to avoid that complication, I'll introduce a new word *very**. This modifier is defined such that if F is vague then *very** F means the same thing as *very* F , and if F is not vague then *very** F is meaningless, like *very forty-seven years old*. It's an excellent question whether *very** is vague, and I think it's a requirement on a definition of vagueness that it allow this question to be asked. As we'll see, this is sadly not true of most proposed definitions of vagueness on the market.

1c. Philosophically Interesting Vague Terms

It's morally obligatory that someone with my standard of living donate 1% of their income to charity. It's not morally obligatory that someone with my standard of living donate 100% of their income to charity. What is the largest x such that it's morally obligatory that someone with my standard of living donate $x\%$ of their income to charity? (As a moralistic cheapskate I'd rather like to know.) Arguably this is vague. But perhaps only arguably. On some divine command

theories it is precise, because there's a fact about what God wants me to do, however hard this is to figure out. (It's even a knowable fact, since God knows it.) But on more standard secular moral theories this may indeed be vague.

There are two lessons to draw from this case. First, if two philosophers can debate what the correct theory of morality is while one thinks it is vague and the other thinks is precise, as I think could happen in a dispute between a divine command theorist and a virtue ethicist, then knowing that a vague term is vague is not required for understanding the term. (I assume here the divine command theorist is not so confused that she's not really talking about *goodness*.)

Second, it is important to remember that for some vague terms competent users of the term need not know in virtue of what they apply. Much of the literature on vagueness focuses on words like *tall*, *thin* and *bald* where all competent users know which kinds of underlying facts are relevant to their application. But not all vague terms are like that, as *good* illustrates. And this phenomena extends beyond the normative, at least narrowly conceived. If you believe Tom Wolfe (2000) then among the youth of America *going out with* is vague and many do not know exactly in virtue of what it applies. It's a familiar point in philosophy of mind that competent users can disagree about what kinds of features a thing must have to satisfy *is thinking*. And we can multiply instances of this by considering any area of philosophy we like.

Id. Discrete Vague Terms

An academic with one child has few children for an academic. An academic with five children does not have few children for an academic. (I'll omit the comparison class 'for an academic' from now on.) Where is the borderline between those with few children and those not with few children? (I don't ask out of personal interest this time.) This question, like the question of how much giving is morally obligatory, feels vague. But note that we cannot generate a compelling Sorites paradox using *has few children*. Let's see how badly this Sorites argument fails.

- (2) a. An academic with one child has few children.
- b. If an academic with one child has few children, then an academic with two children has few children.
- c. If an academic with two children has few children, then an academic with three children has few children.
- d. If an academic with three children has few children, then an academic with four children has few children.
- e. If an academic with four children has few children, then an academic with five children has few children.
- f. So an academic with five children has few children.

Arguably premise e is plausible because as a material conditional it can be seen to be true via the falsity of the antecedent. And at a pinch I can see d as compellingly true for the same reason. But neither b nor c strike me as at all compelling. If someone presents this argument as a Sorites paradox, I simply deny that the paradox-monger can know these premises to be true, or that I have a reason to believe they are true. To be sure, I don't know which premise is false. (If you think you know b to be false replace academics in the example with a more fertile professional group.) But just because I don't know where the argument fails doesn't mean it presents any kind of paradox. When I have no reason to accept two, maybe three, of the premises, the argument falls well short of being paradoxical.

A small note on terminology. Contemporary scientific theories imply that many familiar vague predicates apply in virtue of facts about the world that are, at some level, discrete. What I'm interested in under this heading are predicates where the differences between salient adjacent cases are easily observable, such as the difference between having two and three children.

I.e. Vagueness without Boundaries

The letter of Patrick Greenough's proposal (to be discussed in section three below) suggests that every vague term has only vague boundaries. This is not true. The predicate *in one's early thirties* has a sharp boundary at the lower end and a precise boundary at the upper end. But it isn't too hard to amend his theory to allow for such cases, by saying (in effect) that a vague term is a term with at least one vague boundary. Nicholas Smith makes basically that move in his paper. But such a move won't work, because some vague predicates don't have boundaries. Indeed, some predicates can be vague even though they are satisfied by every object in the domain. The examples here are a little more complicated than in the rest of the paper, but I think they are important enough to warrant the complexity.

For the next several paragraphs the domain will be adult Australian women, and when I use *tall* I'll mean tall for an adult Australian woman. I don't know enough facts to know where the boundaries are for *tall* in this context, but I'll stipulate that a woman shorter than 170cm is determinately not tall, and a woman taller than 180cm is determinately tall. I claim here neither that I know where these boundaries are nor that I could know where they are. But I assume there are boundaries. I'm making these stipulations because it is easier to follow the examples if I use 170 and 180 rather than variables like y and z . It will become obvious that the particular numbers won't matter, as long as there's separation between them. It also doesn't matter whether we use a semantic or epistemic account of determinacy here. It will matter that we use classical logic at various points (e.g. in assuming there are boundaries), but I think that's perfectly reasonable in this context. (Here I follow the arguments in section 2 of Greenough's paper.)

Consider the class of predicates defined by the following schema.

$tall_x =_{df}$ tall or shorter than x cm

For $x < 170$, $tall_x$ has all the same borderline cases as $tall$, and is presumably vague in anyone's book. For $x > 180$, $tall_x$ determinately applies to everyone in the domain, and for now we'll say that makes it not vague. (Though note it need not determinately determinately apply to everyone in the domain, and we'll see below that might be a reason to group it with the vague predicates.)

When x is between 170 and 180, $tall_x$ has some very odd properties. The borderline cases are those women whose height is between x and 180cm. When x is close to 180, this might be a very small border. While we're assuming classical logic, we can assume that there is a value y such that women taller than y cm are tall and those shorter than y cm are not tall. We need not here assume the value of y is either epistemically or *semantically* determinate. Consider a value of x , say 179, such that $x > y$. (Again it's not a necessary assumption that $179 > y$, but it makes the example easier to understand if I use a particular number.) Now $tall_{179}$ has some interesting properties. It has borderline cases, those women between 179 and 180cm tall. But it is satisfied by every woman, since every woman is either tall or shorter than 179cm. I think the existence of the borderline cases is sufficient to make $tall_{179}$ vague. Note that these cases are quite different to *child**, because at the upper boundary there is no sharp jump from borderline cases to clear cases – the two blur together in just the way borderline cases and clear cases of *tall* blur together, so whatever reasons we had to worry about *child** being vague are not applicable here. Still the 'borderline cases' are mislabelled here for there is no border they fall on. Every woman satisfies the predicate. So no definition of vagueness in terms of having a vague boundary, indeed of having a boundary at all, can work.

One might object here that a definition of vagueness is only meant to apply to words not phrases. But just as we can worry about a possible word *child**, we can worry about a possible atomic word *gish* that means the same thing as $tall_{179}$, so that move won't help here.

We now have enough data on the table. In the next section I argue that treating vagueness as being indeterminacy provides a *satisfactory* treatment of the data. In the third section I argue that none of the live alternatives is so satisfactory. So I conclude, somewhat tentatively, that we should define vagueness as indeterminacy.

2. *Vagueness as Indeterminacy*

Back when I was a supervaluationist, I thought that what it was for a term to be vague was for it to refer to different things on different precisifications. That won't do as a theory-neutral definition, for it presupposes supervaluationism, which is not only a theory but a false theory. But we can capture the essential idea in slightly less loaded language.

I will have to make three possibly controversial assumptions. First, I assume a broadly Montagovian perspective, on which we can talk about the referent of an arbitrary term. That referent might be an object, or a truth value, or a function from objects to truth values, or a more complicated function built out of these. Second, I assume we can sensibly use an expanded

Lagadonian language where objects can be names for themselves, truth-values can be names for themselves, functions from objects to truth-values can be names for themselves, and so on. Third, I assume there is no metaphysical vagueness, so each of these Lagadonian names is not vague.

Those assumptions let us make a first pass at a definition of vagueness, as follows. A term t is vague iff there is some object, truth-value or function l which can serve as its own name such that the following sentence is neither determinately true nor determinately false.

(3) t denotes l .

That delivers the intuitively correct account in four of the five cases we discuss above, all except the cases like *child**. I'll say much more about that case below. But it is in one respect slightly too liberal, and we need to make a small adjustment or two to fix this. Consider a predicate F that is defined over a vague domain, but which is determinately satisfied by every object in the domain. Intuitively it is a partial function, which maps every member of its domain to *true*. And assume for sake of argument that it is determinate that it maps every member of the domain to *true*. (Say, for example, it means *is self-identical* when applied to a member of the domain.) Such a predicate is not, I think, vague. But since it is indeterminate which partial function it denotes, the above theory suggests it is vague. We need to make a small adjustment. To state the corrected theory, we will stipulate that *every* term denotes a function. What were previously thought of as terms denoting constants will be treated as terms denoting constant functions. So instead of a name like *Scott Soames* denoting Scott Soames, we'll take it to denote the function that takes anything whatsoever as input, and returning Scott Soames as output. Given that, our second take at a definition of vagueness is as follows.

t is vague iff $\exists x, y_1, y_2$ such that $y_1 \neq y_2$ and it is indeterminate whether $\exists l$ such that t denotes l and $l(x) = y_1$, and it is indeterminate whether $\exists l$ such that t denotes l and $l(x) = y_2$.

To get a sense of the definition, it helps to translate it back into supervaluational talk, and look at the special case where t is a predicate. Then the definition comes to the claim that there is some object that is in the extension of t on one precisification, and in the anti-extension of t on another, which seems like what was intended.

Arguably even that is not enough of a correction. (I'm indebted in the following three paragraphs to Mark Johnston.) Frequently there are debates in semantics over the appropriate *type* of various terms. For instance, a straightforward account would say that in *She ran yesterday*, *yesterday* modifies the intransitive verb *run*, so it denotes a function of type $\langle\langle e,t \rangle, \langle e,t \rangle\rangle$. But on a Davidsonian semantics, *yesterday* denotes a property of the running event being discussed, so its type is simply $\langle e,t \rangle$. Now it is at least a philosophical possibility that there should be no fact of the matter which of these theories is correct.

There are two things we might say about such a possibility. On the one hand, it doesn't at all seem right that a word should count as vague because it is indeterminate what its type should be. That suggests the above definition needs modification. On the other hand, the above definition doesn't imply t is vague whenever there are two distinct functions that could be the denotation of t ; it must also be the case that these functions have overlapping domains. The most natural cases of syntactic indeterminacy are cases where the two possible denotations are functions of quite different types. That suggests the above definition needs no modification.

I think the case for modification is a little stronger. That's partially because the possibility of type-shifting suggests there's a possibility, perhaps a distant one but a possibility, that the second suggestion could fail. And partially because even if there are no uncontroversial cases of syntactic indeterminacy that will mistakenly be treated as cases of vagueness by this theory, the mere possibility of classifying a case of syntactic indeterminacy as a case of vagueness should be enough to warrant concern. And there is a way to modify the definition that does not look like it will lead to mistakenly ruling out any cases of vagueness that should be ruled in, as follows.

t is vague iff $\exists x, y_1, y_2$ such that $y_1 \neq y_2$ and y_1 is of the same type as y_2 , and it is indeterminate whether $\exists l$ such that t denotes l and $l(x) = y_1$, and it is indeterminate whether $\exists l$ such that t denotes l and $l(x) = y_2$.

That implies that if *yesterday* is indeterminate merely because it is indeterminate what type of function it denotes, it won't count as vague, and that's all to the good. So this is our final definition of vagueness.

Still there's a problem with *child**. Many people have thought that it should not be considered vague for one reason or another. Sometimes this is just asserted as a raw intuition, as in Smith and Eklund. There's no arguing with an intuition, so I won't try arguing with it. Rather I'll just repeat a point I made at the start. We aren't here in the business simply of summarising ordinary or philosophical intuitions. Rather we are looking for a definition that captures all the cases that fall into the most theoretically important categories. And intuitions about theoretical importance are less impressive than demonstrations of theoretical importance.

Patrick Greenough (2003) suggests that the problem with terms like *child** is that they aren't vague, but rather that they are simply undefined for the alleged borderline cases. If that's true, and perhaps for some of the examples people had in mind in this area it is, then our definition agrees that they are not vague. For a term that carves a precise division out of part of the domain, and then stays silent, is precise not vague on my account.

Greenough also suggests that the problem with *child** is that it is not higher-order vague. But as he says this can hardly be the entirety of the problem. For it does not seem to be *definitional* that the vague terms are also higher-order vague. True, there is a theoretically

important category of terms that are vague and higher-order vague. But it is not a category that we cannot represent. A term t is in this category just in case t is vague, and *definitely* t is vague, and *definitely definitely* t is vague, and so on. So we can capture that category, even if we don't call only members of that category the vague terms. And this doesn't seem to diminish the theoretical importance of the category of terms I called vague.

It might be thought that what is wrong with *child** is that it cannot be used to generate a Sorites argument. If you think that's what is centrally important to vague terms, then there's a theoretical reason to separate *child** from the genuinely vague. But we should have seen enough by now to show that that can't be right. It's hard to know what it is for a predicate modifier to be Sorites-susceptible, and our last two example predicates, *has few children* and *tall*₁₇₉ cannot be used to set up Sorites arguments. So that *child** does not generate a Sorites paradox is no reason to classify it outside the vague.

So I take it there is no compelling reason to classify *child** and similar terms as precise rather than vague. Admittedly there is an intuition that they are not vague, and perhaps that should be respected. But if the cost of respecting that intuition is that we misclassify several other terms, we should reject the intuition. That's what I'll argue in the next section.

3. Rival Definitions

I just mentioned the idea that a vague predicate could be defined as one that is susceptible to a Sorites argument. This account is sometimes attributed to Delia Graff Fara (2000), but it seems quite a widespread view. For instance, Terence Horgan (1995) says that it is distinctive of vague predicates that they can be used to generate *inconsistency* because the Sorites premises attaching to them are *true*. As I mentioned, such views are vulnerable to a wide variety of counterexamples. Many of these counterexamples also apply to rival definitions of vagueness.

Matti Eklund (2005) develops a similar kind of definition. He starts with Crispin Wright's (1975) famous definition of what it is for a predicate F to be *tolerant*.

Whereas large enough differences in F 's parameter of application sometimes matter to the justice with which it is applied, some small enough difference never thus matters.

Eklund's position then is that F is vague iff it is part of semantic competence with respect to F to be disposed to accept that F is tolerant. Eklund agrees that it is inconsistent to assert that F is indeed tolerant. But as he has argued extensively elsewhere, the falsity of the tolerance principle is compatible with it being part of competence that one is disposed to accept it. (A view in the same family is put forward in Sorensen 2001.) I have no wish to dispute this part of Eklund's theory. Indeed that meaning principles can be false, even inconsistent, it seems to have been a fairly fruitful idea in a variety of areas of Eklund's philosophy. But I don't think it helps with vague terms.

Three of the problems with this have already been given. It is not clear what a parameter of application for a non-predicate like *very* even is, so it isn't clear what it means to say that *very* is tolerant. It surely is not required of competent users of *few children* that they are disposed to accept the premises in our earlier Sorites argument. And for some vague predicates, like *tall*₁₇₉, the tolerance principle is not plausible to a competent speaker because it is not plausible that a "large enough" difference in the parameter of application (presumably height) matters. These problems all seem to carry over from the problems associated with Sorites based definitions.

I suspect, though I'm less certain here, that the philosophically interesting cases also pose a problem for Eklund's view. When we look at philosophically interesting cases, like *being good*, there are two distinct ways to read Eklund's claim that competent speakers are disposed to accept the tolerance principle. These are the wide scope and the narrow scope reading. To see the ambiguity, let's write out Eklund's principle in full.

Competent speakers are disposed to accept that whereas large enough differences in *F*'s parameter of application sometimes matter to the justice with which it is applied, some small enough difference never thus matters.

Here's the wide scope reading of this.

F's parameter of application is such that whereas competent speakers are disposed to accept that large enough differences in it sometimes matter to the justice with which *F* is applied, some small enough difference never thus matters.

And here is the narrow scope reading, with a phrase added for emphasis.

Competent speakers are disposed to accept that whereas large enough differences in *F*'s parameter of application, *whatever it is*, sometimes matter to the justice with which it is applied, some small enough difference never thus matters.

To see the difference between the two cases, assume for the sake of argument that a competent speaker thinks that to be good is to do actions whose consequences have a high enough utility, whereas in reality to be good is to obey enough of God's commands. In each case *being good* is vague, because we are using satisficing versions of consequentialism and divine command theory. So the parameter of application for *being good* is the number of God's commands you obey. The competent speaker will not accept the wide scope version of tolerance with respect to *being good*, because they don't think that large differences with respect to how many of God's commands you obey matter to the justice with which *being good* is applied. Such cases can be multiplied endlessly to show that the wide scope version of Eklund's principle cannot generally be true,

because it makes it the case that competent speakers have correct views on contentious philosophical matters the resolution of which goes beyond semantic competence. For these reasons Eklund has said (personal communication) that he intends the narrow scope version.

But the narrow scope version also faces some difficulties. The most direct problem is that one can be a competent user of a term like *food* or *dangerous* or *beautiful* without having any thoughts about parameters of application. I suspect I was a competent user of these terms before I even had the concept of a parameter of application. Even bracketing this concern, there is a worry that competence requires knowing of a term whether it is vague or not. But this seems to be a mistake. It is not a requirement of competence with moral terms like *good* that one know whether they are maximising or satisficing terms. Tom Wolfe and the students he observed while writing *I Am Charlotte Simmons* seemed to disagree about whether *going out with* is vague, but they were both competent users, they simply disagreed on something like a normative question. (See Wolfe 2000 for more on his take on matters.) And it seems that two users of language could disagree over whether *is thinking* is vague without disagreeing over whether either is a competent semanticist. They may well disagree over whether either is a competent philosopher of mind, but such disagreements are neither here nor there with respect to our present purposes. So I don't think that either disambiguation of Eklund's principle can properly account for vagueness in philosophically interesting terms.

Nicholas Smith argues for a definition of vagueness that uses some heavier duty assumptions about the foundations of semantics. In particular, he sets out the following definition,

Closeness If *a* and *b* are very similar in *F*-relevant respects, then '*Fa*' and '*Fb*' are very similar in respect of truth.

and goes on to say that vague predicates are those that satisfy non-vacuously satisfy Closeness over some part of their domain. For this to work there must be, as Smith acknowledges, both degrees of truth and something like a distance metric defined on them. (These are separate assumptions; in the theory of Weatherston 2005 the first is true but not the second.) I won't question those assumptions, but rather focus on the problems the definition has even granting the assumptions.

As with the two definitions considered so far, it is hard to see how this could possibly be generalised to cover vagueness in non-predicates. It's true (given our assumptions) that if *a* and *b* are similar in *very tall*-relevant respects, then '*a* is very tall' and '*b* is very tall' will be similar in respect of truth. But that doesn't show *very* is vague, for the same condition is satisfied when we replace *very* with the precise modifier *doubly*. This isn't an argument that Smith's definition *couldn't* be extended to cover modifiers, but a claim that it is hard to see how this will work.

The definition also has trouble with *tall*₁₇₉ for this satisfies Closeness vacuously. Though to be fair given the logical assumptions Smith makes, it is possible that no predicate with the properties I've associated with *tall*₁₇₉ can be defined.

More seriously, there is a problem with predicates like *has few children*. It just isn't true that "An academic with two children has few children" is close in truth value to "An academic with one child has few children". In general Smith's theory has trouble with, i.e. rules out by definition, vague terms where the underlying 'relevant respects' are highly discrete. Note that the problem here extends to some predicates where the underlying facts are continuous. Consider the predicate *is very late for the meeting*. At least where I come from, a person who is roughly ten minutes late is a borderline case of this predicate. But which side of ten minutes late they are matters. (In what follows I make some wild guesses about how numerical degrees of truth, which aren't part of my preferred theory, should operate. But I think the guesses are defensible given the empirical data.) If Alice is nine and three-quarters minutes late, and Bob is ten and a quarter minutes late, then the degree of truth of "Alice is very late" will be much smaller than the degree of truth of "Bob is very late". The later you are the truer "you are very late" gets, but crossing conventionally salient barriers like the ten minutes barrier matter much more to the degree of truth than crossing other barriers like the nine minutes thirty-three seconds barrier. Smith (in conversation) has suggested that he's prepared to accept that *is very late for the meeting* is only partially vague if the truth-values 'jump' at the ten minute mark as I'm suggesting. But this seems improper, for this is as clear a case of a vague predicate as we have. Still, it's worth remembering as always that every definition has its costs, and this may be a cost one chooses to live with. Personally I think it is excessive.

Patrick Greenough did not put forward his theory as a *definition* of vagueness, but rather as a *minimal theory* to which all partisans could agree. Like Eklund, Greenough plays off Crispin Wright's idea of tolerance. Roughly, a vague predicate is one that is epistemically tolerant – it's one where you can't know that a small difference makes a difference. Here's a less rough statement of it, though note this is heavily paraphrased.

Let τ be a variable that ranges over truth states (e.g. true, determinately true, not determinately determinately not determinately true, etc.) v a function from objects to real numbers such that whether x is F depends only on the value of $v(x)$ (i.e. v is F 's parameter of application) and c a suitably small number. Then F is vague iff the following claim non-vacuously holds.

$\forall \tau \forall \alpha \forall \beta \forall a \forall b$, if $|v(\alpha) - v(\beta)| < c$ and a names α and b names β and it is knowable that Fa is τ then it is not knowable that Fb is not τ .

Less formally, we can't know where any boundary at any order of definiteness for F lies. (It isn't clear in Greenough's presentation exactly what the non-vacuous condition comes to. He only

explicitly says that for the special case where τ is *is true* there must be an a and a b such it is knowable that Fa is τ and Fb is not τ , but maybe that should be extended to all τ .) Because of cases like *in one's early thirties* this cannot do as a general definition, but it is easy enough to repair it by restricting the quantifier attaching to a and b to a range over which F has only vague boundaries. Doing this amounts to weakening Greenough's claim from the view that vague terms have only vague boundaries to the view that they have some vague boundaries, which seems plausible. But still there are problems.

Most obviously, *tall*₁₇₉ does not non-vacuously satisfy the tolerance requirement. And like all the tolerance-based theories it is far from clear how it should be extended to vagueness in non-predicates. On the other hand, Greenough's theory might well handle the discrete cases like *has few children*. I say *might* rather than *does* because it is rather hard to work out how the higher-orders of vagueness work for such terms. I'll simply note that there are some plausible enough epistemic models on which *has few children* satisfies his requirement.

There is a problem which is distinctive to Greenough's view of his theory as a minimal theory. As Smith notes, Greenough makes it a requirement that vague boundaries are unknown. But this is controverted in some mainstream theories, for example the version of supervaluationism in Dorr (2003). Since Dorr's theory should not be ruled out by a minimal theory or a definition, this is a weakness in Greenough's theory.

The more philosophically interesting problems concern, appropriately enough, the philosophically interesting terms. Greenough has a proof that his definition is equivalent to a definition in terms of borderline cases. The proof has several assumptions, one of which being that we know what the parameter of application of a vague term is. More precisely, he assumes that we know everyone older than an old person is old, which is unproblematic, but he also assumes that the proof generalises to all vague cases, and this amounts to the assumption that we know parameters of application. As we've seen, this isn't true of philosophically interesting vague terms. This leaves open the possibility that Greenough's theory, unlike Smith's and Eklund's theories, overgenerates. The following is probably not a live possibility in any interesting sense, but it isn't I think the kind of thing a definition (or minimal theory) should rule out by definition.

It is possible that a kind of mysterianism about ethics is true, and we cannot know whether *good* is vague or precise. For a concrete example, let's assume it is knowable that some kind of divine command theory is true, but it is unknowable whether to be good one must obey all of God's commands or merely enough of them, where it is vague what counts as enough of them. In fact morality requires obeying all God's commands, but this is not knowable – for all we know the satisficing version is the true moral theory. If this is the case then *good* will be epistemically tolerant, for we cannot know that a small difference in how many of God's commands you obey makes a difference to whether you are good, or determinately good etc. But in fact *good* is precise, for it precisely means obeying all of God's commands. Earlier I objected to Eklund's

theory because semantic competence does not require knowing parameters of application, especially as such. This is the converse objection – I claim that a term's being precise does not imply that we know, or even *could know*, that it applies in virtue of a precise condition. All that matters is that it *does* apply in virtue of a precise condition.

It's a constant danger in philosophy that one infer from the falsity of all extant rivals that one's preferred theory is correct. I certainly don't want to argue that because Eklund's, Smith's and Greenough's definitions are incorrect that the traditionalist definition I have offered must be right. But we can make that conclusion more plausible by noting how widely the arguments levelled here generalise. The philosophically interesting cases seem to tell against *any* definition of vagueness in terms of semantic competence, for they show that competent users can have exactly the same attitude towards vague terms as they have towards precise terms. And our moral example suggests that any definition in terms of epistemic properties will be in trouble for it might not be knowable whether a particular term is vague or precise. Finally, the cases of vague predicate modifiers raise difficulties for any attempt to define the vagueness of a term in terms of properties of sentences in which it is used rather than mentioned. For it seems that as long as *very** attaches only to vague predicates, then whether *very** is vague or precise will make no salient differences to the sentences in which it appears. So we have to look at sentences in which the allegedly vague term is mentioned. And while I don't have a definitive argument here, I think looking at the range of cases we want to cover, and in particular at the range of cases where tolerance-type principles fail to be non-vacuously satisfied, our best option for completing these sentences is to look whether the term has a determinate or indeterminate denotation. We can then pass the questions of what determinacy consists in, and in particular the question of whether it is an epistemic or semantic feature, to the theorist of vagueness.

References

- Dorr, Cian (2003) "Vagueness without Ignorance" *Philosophical Perspectives* 17: 83-113.
- Eklund, Matti (2005) "What Vagueness Consists in" *Philosophical Studies* 125: 27-60.
- Field, Hartry (1973) "Theory Change and the Indeterminacy of Reference" *Journal of Philosophy* 70: 462-81.
- Fine, Kit (1975) "Vagueness, Truth and Logic" *Synthese* 30: 265-300.
- Fara, Delia Graff (2000) "Shifting Sands: An Interest-Relative Theory of Vagueness" *Philosophical Topics* 28: 45-81. [This paper was published under the name "Delia Graff".]
- Greenough, Patrick (2003) "Vagueness: A Minimal Theory" *Mind* 112: 235-81.
- Horgan, Terrence (1995) "Transvaluationism: a Dionysian Approach to Vagueness" *Southern Journal of Philosophy* 33, Spindel Conference Supplement, 97-125.
- Sainsbury, Mark (1991) "Is There Higher-Order Vagueness?" *Philosophical Quarterly*, 41: 167-82.
- Smith, Nicholas J. J. (2005) "Vagueness as Closeness" *Australasian Journal of Philosophy* 83: 157-83.
- Sorensen, Roy (2001) *Vagueness and Contradiction*, Oxford: Oxford University Press.
- Weatherson, Brian (2005) "True, Truer, Truest" *Philosophical Studies* 123: 47-70.
- Wolfe, Tom (2000) *Hooking Up*, New York: Farrar, Strauss and Giroux.
- Wright, Crispin (1975) "On the Coherence of Vague Predicates" *Synthese* 30: 325-65.