

9 Let us take stock. Vagueness should be characterized as boundarylessness, not merely in terms of borderlines. Boundarylessness cannot be described sharply, for example set-theoretically; so, whatever insight psychological descriptions may offer, the only semantic description which appears plausible is vague, for example homophonic. We must reject the classical picture of classification by pigeon-holes, and think in other terms: classifying can be, and often is, clustering round paradigms.

Just how widespread vagueness is can be underestimated. Let me draw attention to an area sometimes wrongly thought to be free of it: biological species. Even in quite recent philosophy, there is a tendency to suppose that species come in the "eternal and fixed forms" beloved, according to John Locke, of the Port Royal logicians. It may seem that *strawberry* draws boundaries, since there are no borderline cases. But this is just an accident. There could very well be, and no doubt with the advent of genetic engineering soon will be, a series of plants between strawberries and raspberries, many of them borderline for both concepts. Such concepts do not impose boundaries, but constitute one of the largest and most impressive systems of contrary boundaryless concepts. Locke was right to draw attention to the lack of boundaries by reminding us of boundary-defying "monsters".

One practical application of work on vagueness is in cognitive science, where a possible goal is to implement in machinery the vagueness of our concepts. Another application has already been mentioned. The law must rule a boundary between legitimate and illegitimate acts. Here, boundarylessness would be out of place. Yet such rulings must often traverse territory spanned by a boundaryless concept, like that of being a person. Given the nature of boundarylessness, semantics give freedom. There is some number of minutes such that the nature of the concept of a person, together with the nature of the world, makes it neither mandatory nor impermissible to apply the concept to a foetus of that age in minutes. Hence arguments that use the vague concept to establish or overthrow a sharp ruling are alike inadequate. We can no more argue that aborting a foetus of this age is right because it is not a person than we can argue that it is wrong because it is a person, if *person* is vague at the crucial point. In general, only a pragmatic justification could be found for drawing a legal line in an area where there are no relevant boundaries.

I mention this merely as an example of a possible application whose details remain to be worked out. It proleptically exemplifies my hope that work in the philosophy of vagueness will enable us better to understand how the demands of law and morality should be tailored to the boundaryless fabric of most of our thought and talk.

14 Vagueness and ignorance

Timothy Williamson

No one knows whether I'm thin. I'm not clearly thin; I'm not clearly not thin. The word "thin" is too vague to enable "TW is thin" to be recognized as true or as false, however accurately my waist is measured and the result compared with vital statistics for the rest of the population. Is this ignorance? Most work on vagueness has taken for granted the answer "No". According to it, there is nothing here to be known. I am just a borderline case of thinness; "TW is thin" is neither true nor false. Doubt will be cast on the coherence of this view. There are standard objections to the alternative that "TW is thin" is either unknowably true or unknowably false. Doubt will be cast on them too. For all we know, vagueness is a kind of ignorance.

1 Why doubt the majority view? Well, suppose that "TW is thin" is neither true nor false. If I were thin, "TW is thin" would be true; since it isn't, I'm not. But if I'm not thin, "TW is not thin" is true, and so "TW is thin" false. The supposition seems to contradict itself. Yet on the majority view it is true.

To generalize the argument, consider a language L with negation (\sim), disjunction (\vee), conjunction ($\&$) and a biconditional (\leftrightarrow). Extend L to a metalanguage for L by adding a truth predicate (T) for sentences of L and quotation marks (" \dots ") for naming them. The falsity of a sentence of L is identified with the truth of its negation. Thus the supposition at issue, the denial of bivalence for a sentence of L, is equivalent to the denial that either it or its negation is true:

$$(1) \quad \sim[T("P") \vee T("\sim P")]$$

Two instances of Tarski's disquotational schema for truth are:

$$(2a) \quad T("P") \leftrightarrow P$$

$$(2b) \quad T("\sim P") \leftrightarrow \sim P$$

The argument uses (2a) and (2b) to substitute their right-hand sides for their left-hand sides in (1):

$$(3) \quad \sim[P \vee \sim P]$$

It then applies one of De Morgan's laws to (3), giving

$$(4) \quad \sim P \ \& \ \sim \sim P$$

This is a contradiction, whether or not the double negation is eliminated. Thus (1) reduces to absurdity. In effect, one uses Tarski's schema to equate bivalence ($T("P") \vee T("\sim P")$) with the law of excluded middle ($P \vee \sim P$), and then argues from the incoherence of denying the latter to the incoherence of denying the former.¹

The argument does not purport to show that bivalence must be asserted, only that it must not be denied. Whether bivalence must be asserted will depend on whether the law of excluded middle must be asserted, an issue which has not been addressed. Even so, the argument may seem to prove too much. Can every denial of bivalence be reduced to absurdity? However, the argument applies not whenever bivalence is denied, but only when it is denied of a particular sentence. It does not touch intuitionism in mathematics, for example. Although intuitionists deny the general principle of bivalence, they are forbidden to give particular counterexamples, just because the inference from (1) to (4) is intuitionistically valid.² They sometimes refrain from asserting the bivalence of a particular sentence, but they never deny it. This does not undermine their denial of the general principle, for "Not every sentence is bivalent" does not intuitionistically entail "Some sentence is not bivalent". Vagueness is a different matter. Vague sentences are supposed to be obviously not bivalent in borderline cases, and the usual way of evoking this sense of obviousness is by vivid descriptions of particular examples. If it is obvious that not all vague sentences are bivalent, it is obvious that "TW is thin" is not bivalent. So if one must not deny the bivalence of "TW is thin", does vagueness give any reason to deny bivalence in general?³

The core of the argument is its use of Tarski's disquotational schema for truth; everything else is relatively uncontroversial.⁴ At first sight, it

1. Tarski derives bivalence (which he calls "the principle of excluded middle") from his definition of truth (rather than the disquotational schema) as Theorem 2 of his 1931 (see Tarski 1956, p. 197). The proof uses the law of excluded middle (in the present sense).
2. The intuitionist is assumed here to equate "true" with "provable" rather than with "proved".
3. The sorites paradox might move some to deny that all sentences of the form "*n* grains make a heap" are bivalent, but not to deny bivalence for any particular *n* (although for some *n* they would refrain from asserting it). If the argument in the text is sound, the supposedly obvious assumptions which might drive one to this view are false.
4. A general setting for the argument is as follows. There is a partial ordering \leq of the semantic values assigned to sentences (e.g. truth values) under which they form

looks vulnerable to an obvious objection. If *P* is neither true nor false, should not $T("P")$ be simply false? But then the left-hand sides of (2a) and (2b) but not their right-hand sides would be false; on a strong reading of the biconditional this would make (2a) and (2b) not true. This is just what one might say in the case of reference failure. Consider, for example, a context where "this dagger" fails to pick anything out. One might hold that no sentence in which "this dagger" is used is true in the context. Thus neither "This dagger is sharp" nor "This dagger is not sharp" is true, so "This dagger is sharp" is neither true nor false. By the same principle, the Tarskian biconditional "This dagger is sharp" is true if and only if this dagger is sharp" would not be true in the context, for it uses as well as mentions "this dagger". The argument has no force in the case of reference failure; why should it have any force in the case of vagueness?

On the suggested treatment of reference failure, "This dagger is sharp" says nothing that could have been true or false, and even counterfactuals such as "If the servants had been assiduous, this dagger would have been sharp" are neither true nor false (of course, the sentence type "This dagger is sharp" could have been used in a different context to say something true or false). According to the parallel treatment of a vague sentence in a borderline case, "TW is thin" says nothing that could have been true or false, and even counterfactuals such as "If he had dieted, TW would have been thin" are neither true nor false.⁵ Such consequences are unwelcome. Unlike "This dagger is sharp", "TW is thin" could have said something true without saying

a lattice, i.e. each pair of values has a greatest lower bound (glb) and a least upper bound (lub), greater values being thought of as "truer". If $|P|$ is the semantic value assigned to *P*, $|P \ \& \ Q| = \text{glb}(|P|, |Q|)$, $|P \ \vee \ Q| = \text{lub}(|P|, |Q|)$ and if $|P| \leq |Q|$ then $|P \ \vee \ Q| \leq |P|$. These assumptions are met by standard classical, supervaluational, intuitionist and many-valued treatments and others. Now suppose that $|P| = |T("P")|$ and $|\sim P| = |T("\sim P")|$. Then $|P| = |T("P")| \leq \text{lub}(|T("P")|, |T("\sim P")|) = |T("P") \vee T("\sim P")|$; similarly, $|T("\sim P")| \leq |T("P") \vee T("\sim P")|$. Thus $|\sim [T("P") \vee T("\sim P")]| \leq |\sim P|$ and $|\sim [T("P") \vee T("\sim P")]| \leq |\sim P|$, so $|\Omega| \leq \text{glb}(|\sim P|, |\sim \sim P|) = |\Omega|$. Thus what is needed is a defence of Tarski's schema which assigns the same semantic value to each side of the biconditional; this is supplied in the text. Two further assumptions are that *P* has a negation whose falsity is equivalent to the truth of *P* and that a contradiction is indeed absurd. The former is clearly correct for "TW is thin", which is enough for the argument. As for the latter, if the denial of bivalence for vague sentences is obviously correct, it does not involve a contradiction (someone might answer the question "Is TW thin?" with "He is and he isn't", but it would take a bold man to revise logic on the basis of that idiom).

5. A standard analysis of "If he had dieted, TW would have been thin" is assumed, on which it results from feeding "He [or TW] dieted" and "TW is thin" into a counterfactual conditional.

something different. Most simply, "TW is thin" means that TW is thin; on the suggested treatment, "This dagger is sharp" means that this dagger is sharp" is neither true nor false, for "this dagger" is used, not mentioned, on its second occurrence. But since "TW is thin" means that TW is thin, what it is for "TW is thin" to be true is just for TW to be thin. Similarly, since "TW is not thin" means that TW is not thin, what it is for "TW is not thin" to be true, and so for "TW is thin" to be false, is just for TW not to be thin.⁶ The difference between reference failure (as treated above) and vagueness favours the Tarskian biconditionals in the latter case.⁷ In doing so, it undermines the thought that "TW is thin" is neither true nor false, i.e. (1), by vindicating the argument from it to (4).

We can consistently deny bivalence of a sentence with reference failure precisely because in doing so we abjure the use, embedded or unembedded, of that sentence in that context. If we are not willing to abjure the use, embedded or unembedded, of a vague sentence in the context of borderline cases, we cannot consistently deny its bivalence. According to a sceptical view, rigour demands that we should abjure such uses because vagueness is itself a kind of reference failure. Adjectives refer, if at all, to sharply defined properties, but a vague one like "thin" fails to single out such a property and so fails to refer; sentences of the form "a is thin" say, strictly, nothing, whether or not *a* is a borderline case. Since almost all our utterances involve vague terms, this view

6. A homophonic truth theory is thus not essential to the argument; the translatability of *P* and $\sim P$ into the metalanguage is enough, as Tarski noted. Incidentally, it is not claimed that a Tarskian theory tells the whole truth about truth, just that it tells an essential part of the truth. Without a disquotational schema, it is doubtful that one has a truth predicate at all.

7. The supervenient treatment of vagueness, most systematically expounded by Kit Fine (this volume), may seem an obvious counterexample to the argument in the text from failure of bivalence to failure of excluded middle. However, Fine allows a Tarskian truth predicate "true_T"; he argues that it is conceptually prior to the ordinary truth predicate "true", because "x is true" is to be defined by "Definitely (x is true)_T"; "true" is not subject to the Tarskian schema (pp. 148–49; compare Kripke 1975, p. 715). Since Fine's account validates the law of excluded middle, it validates bivalence for the primary notions of truth and falsity. Where the present approach differs is in its claim that the ordinary notion of truth is subject to the Tarskian schema and is therefore not to be defined in Fine's way. The "definitely" operator is discussed in section 2. In *The Logical Basis of Metaphysics*, Michael Dummett argues that the ordinary notion of truth for a vague language is the non-Tarskian one because only it is "objective" in the sense that "every sentence determinately either does or does not possess it" (Dummett 1991, p. 74). This condition is not self-evident; not least because it is inconsistent with second-order vagueness. The disquotational schema is endorsed for a vague language by MacIntyre (this volume, p. 201); Peacocke 1981, pp. 136–7 (both within theories of degrees of truth); Sainsbury (this volume).

makes almost all of them mere noise. They are not even failed attempts to express thoughts, since parallel considerations would suggest that almost all our concepts are equally contentless.⁸ The only consistent expression of such a view is in silence.⁹

Once we are permitted to use "thin", we can argue that "TW is thin" says something that would have been true in various circumstances, because I would have been thin. Then "TW is thin" is true if and only if TW is thin" says something too. But if it says anything, it is true. For, given that "TW is thin" means that TW is thin, what more could it take for "TW is thin" to be true than for TW to be thin?

To deny bivalence for vague sentences while continuing to use them is to adopt an unstable position. The denial of bivalence amounts to a rejection of the practice of using them. One is rejecting the practice while continuing to engage in it. Rapid alternation between perspectives inside and outside the practice can disguise, but not avoid, this hypocrisy.

2 If one cannot deny bivalence for vague sentences, can one deny something like it? There is a standard move at this point. Instead of saying that "TW is thin" is neither true nor false, one says that it is neither *definitely* true nor *definitely* false. Definite truth does not itself obey the disquotational schema, otherwise nothing would have been gained. It takes less for "TW is thin" not to be definitely true than for TW not to be thin. Since it does not take less for "TW is thin" not to be true than for TW not to be thin, truth is not the same thing as definite truth. On pain of the argument in section 1, this new position does not involve a denial of bivalence. Indeed, the principle of bivalence does not mention definiteness; it merely says that a sentence is either true or false. On the face of it, the claim that a sentence is neither definitely true nor definitely false has no more to do with bivalence than the claim that it is neither necessarily true nor necessarily false, or that it is neither obviously true nor obviously false.¹⁰ To pursue indirect connections would be premature.

Before one can assess the claim that vague sentences are neither definitely true nor definitely false in borderline cases, one needs to know what it means. That the adverb "definitely" has been given a clear

8. Even the intention to express some thought or other harbours vagueness.

9. The classic expression of a sceptical view is section 56 of Frege 1903. The more limited sceptical view that observational predicates are vague in such a way as to be incoherent is discussed in my 1990a, pp. 88–103; in effect section 6 below explains how the incoherent principle "if *x* and *y* are indiscriminable by the naked eye, *x* is thin if and only if *y* is thin" could look true while being false.

10. For a contrary view see Dummett 1991, pp. 74–82.

relevant sense is less than obvious. If "definitely true" were just a circumlocution for "true", no problem would arise, but the view under consideration requires the two expressions to have quite different senses. Can "definitely" be explained in other terms, or are we supposed to grasp it as primitive? No doubt "TW is thin" is definitely true if and only if TW is definitely thin, but what is the difference between being thin and being definitely thin? Is it like the difference between being thin and being very thin? Again, "TW is thin" is presumably not definitely true if and only if TW is not definitely thin; what is the difference between not being thin and not being definitely thin?¹¹

Let it be obvious that "TW is thin" is neither definitely true nor definitely false. In reporting this obvious truth, the philosopher has no right to stipulate a theoretical sense for "definitely". Rather, it must be used in a sense expressive of what is obvious. Yet what is *obvious* is just that vague sentences are sometimes neither knowably true nor knowably false. The simplest hypothesis is that this is the *only* sense in which the vague sentences are neither definitely true nor definitely false. Bivalence and classical logic hold. Either I'm thin and "TW is thin" is true or I'm not thin and "TW is thin" is false; we have no way of knowing which. Although this is not at all the standard view of what "definitely" means, the obscurity of the standard view gives us reason to explore alternatives. The epistemic view is usually held to be inconsistent with obvious facts, but the leading candidate for such a fact—the failure of bivalence—has already disappeared. The rest of the paper explores the epistemic view.¹²

3 Many descriptions of vagueness rule out the epistemic view from the start. A term is said to be vague only if it can have a borderline case, and a case is said to be borderline only if our inability to decide it does not depend on ignorance. But to assume that the cases ordinarily called "borderline" are borderline in this technical sense is just to beg the

11. There are views on which "definitely" makes a difference only in the scope of negation and in similar contexts. I also assume that the reference of "TW" is unproblematic.

12. The epistemic view probably goes back to the Stoic logician Chrysippus, a man with some claim to have discovered the classical propositional calculus; see Barnes 1982 and Burnyeat 1982. More recently, it has been defended in Cargile (this volume); Campbell 1974; Sorensen 1988, pp. 217–52; Horwich 1990, pp. 81–87; Williamson 1990a, pp. 103–8. It is critically discussed in Heller 1990, pp. 89–106. Sperber and Wilson 1986 retains classical logic and semantics while explaining vague utterances in pragmatic terms, "He is bald", said of a man with just one hair on his head, being false but relevantly informative; a generalization of this view from "bald" to "heap" might need to postulate ignorance.

question against the epistemic view. For example, "TW is thin" would ordinarily be called a "borderline" case, but one should not assume without argument that our inability to decide the matter does not depend on ignorance. Of what fact could we be ignorant? There is an obvious answer: we are ignorant either of the fact that TW is thin or of the fact that TW is not thin (our ignorance prevents us from knowing which). If that is a bad answer, it has yet to be explained why. That it uses the word "thin" is just what one would expect in the light of section 1. There is no general requirement that vague words be definable in other terms.

Those wholly predictable opening moves against the epistemic view mismanage a deeper objection. It can be made using the idea that vague facts *supervene* on precise ones. If two possible situations are identical in all precise respects, they are identical in all vague respects too. For example, if x and y have exactly the same physical measurements, x is thin if and only if y is thin. More generally:

- (*) If x has exactly the same physical measurements in a possible situation s as y has in a possible situation t , x is thin in s if and only if y is thin in t .¹³

The objection to the epistemic view can now be formulated. Let my exact physical measurements be m . According to the epistemic view, I am either thin or not thin. By (*), if I am thin, necessarily anyone with physical measurements m is thin. Similarly, if I am not thin, necessarily no one with physical measurements m is thin. Thus either being thin is a necessary consequence of having physical measurements m , or not being thin is. Suppose that I find out, as I can, what my physical measurements are. I would then seem to be in a position either to deduce that I am thin or to deduce that I am not thin. But it has already been conceded that no amount of measuring will enable me to decide whether I am thin.¹⁴

The basis of this objection to the epistemic view is not that one can know all the relevant facts in a case ordinarily classified as "borderline" but that one can know a set of facts on which all the relevant facts supervene, without being able to decide the case. Unlike the first claim, the second does not beg the question against the epistemic view. The epistemic theorist has as much reason as anyone else to accept supervenience claims like (*). However, the objection commits a subtler fallacy.

13. More accurately, one's thinness may depend on the physical measurements of one's comparison class as well as on one's own. This does not affect the point about to be made.

14. Exercise: how does this argument fare against the supervenience approach?

The kind of possibility and necessity at issue in supervenience claims like (*) is metaphysical. There *could not* be two situations differing vaguely but not precisely. Suppose that I am in fact thin. By (*), it is metaphysically necessary that anyone with physical measurements *m* is thin. If I know that I have physical measurements *m*, in order to deduce that I am thin I must *know* that anyone with physical measurements *m* is thin. The plausibility of the objection to the epistemic view thus depends on something like the inference that since the supervenience generalizations are metaphysically necessary, they can be known *a priori*. The inference from metaphysical necessity to *a priori* knowability may be a tempting one: but, as Kripke has emphasized, it is fallacious. Indeed, metaphysical necessities cannot be assumed to be knowable in any way at all, otherwise all mathematical truths could be assumed knowable. It is integral to the epistemic view that metaphysically necessary claims like "Anyone with physical measurements *m* is thin" can be as unknowable as physically contingent ones like "TW is thin".

One should not be surprised that the known supervenience of *A*-facts on *B*-facts does not provide a route from knowledge of *B*-facts to knowledge of *A*-facts. A more familiar case is the supervenience of mental facts on physical facts. Suppose, for the sake of illustration, that bravery is known to supervene on the state of the brain. Then if *s* is a maximally specific brain state (described in physical terms) of brave Jones, it is metaphysically necessary that anyone in brain state *s* is brave. Clearly, however, there is no presumption that one could have found out that Jones was brave simply by measuring his brain state and invoking supervenience. "Anyone in brain state *s* is brave" cannot be known *a priori*. Perhaps one can know it *a posteriori*, because one can find out that someone is brave by observing his behaviour, then combine this knowledge with knowledge of his brain state and of the supervenience of mental states on brain states. "Anyone with physical measurements *m* is thin" cannot be known *a posteriori* in a parallel way, for no route to independent knowledge of someone with physical measurements *m* that he is thin corresponds to the observation of brave behaviour.

The epistemic view of vagueness is consistent with the supervenience of vague facts on precise ones. The next section considers a different objection to the epistemic view, and makes another application of the concept of supervenience.

4 A common complaint against the epistemic view of vagueness is that it severs a necessary connection between meaning and use. Words mean what they do because we use them as we do; to postulate a fact of the matter in borderline cases is to suppose, incoherently, that the

meanings of our words draw lines where our use of them does not. The point is perhaps better put at the level of complete speech acts, in terms of sentences rather than single words. The meaning of a declarative sentence may provisionally be identified with its truth conditions, and its use with our dispositions to assent to and dissent from it. The complaint is that the epistemic view of vagueness sets truth conditions floating unacceptably free of our dispositions to assent and dissent.

So far, the complaint is too general to be convincing. If our dispositions to assent to or to dissent from the sentence "That is water" do not discriminate between H_2O and XYZ, it does not follow that the truth conditions of the sentence are equally indiscriminating. What needs to be emphasized is that there is no sharp natural division for the truth conditions of "He is thin" to follow corresponding to the sharp natural division between H_2O and XYZ followed by the truth conditions of "That is water". The idea is that if nature does not draw a line for us, a line is drawn only if we draw it ourselves, by our use. So there is no line, for our use leaves not a line but a smear.

Before we allow the revised complaint to persuade us, we should probe its conception of drawing a line. On the face of it, "drawing" is just a metaphor for "determining". To say that use determines meaning is just to say that meaning *supervenes* on use. That is: same use entails same meaning, so no difference in meaning without a difference in use. More formally:

(#) If an expression *e* is used in a possible situation *s* in the same way as an expression *f* is used in a possible situation *t*, *e* has the same meaning in *s* as *f* has in *t*.

There are various problems with (#), such as its neglect of the environment as a constitutive factor in meaning and its crude notion of "used in the same way". However, some refinement of (#) will be assumed for the sake of argument to be correct. For the epistemic view of vagueness is quite consistent with (#) and its refinements. Although the view does not permit simple-minded reductions of meaning to use, it in no way entails the possibility of a difference in meaning without any corresponding difference in use. Had "TW is thin" had different truth conditions, our dispositions to assent to and dissent from it would have been different too.

Our use determines many lines. Of these one of the least interesting is the line at which assent becomes more probable than dissent. It is no more plausible a candidate for the line between truth and falsity than is the line at which assent becomes unanimous. The study of vagueness has regrettably served as the last refuge of the consensus theory of truth; the theory is no more tenable for vague sentences than it is for

precise ones. We can be wrong even about whether someone is thin, for we can be wrong both about that person's shape and size and about normal shapes and sizes in the relevant comparison class. These errors may be systematic; some people may characteristically look thinner or less thin than they actually are, and there may be characteristic misconceptions about the prevalence of various shapes and sizes. To invoke perfect information or epistemically ideal situations at this point is merely to swamp normal speakers of English with more measurements and statistics than they can handle. Perhaps an epistemically ideal speaker of English would be an infallible guide to thinness, but then such a speaker might know the truth value of "TW is thin". If one sticks to actual speakers of English, there is no prospect of reducing the truth conditions of vague sentences to the statistics of assent and dissent, whether or not one accepts the epistemic view of vagueness.

The failure of simple-minded reductions is quite consistent with supervenience. There may be a subtler connection, perhaps of a causal kind, between the property of thinness and our use of "thin". Even if everything has or lacks the property, the reliability of our mechanism for recognizing it may depend on its giving neither a positive nor a negative response in marginal cases. The cost of having the mechanism answer in such cases would be many wrong answers. It is safer to have a mechanism that often gives no answer than one that often gives the wrong answer. From such a mechanism, one might be able to work back to the property, through the question "Which property does this mechanism best register?".¹⁵

It might be objected that if a mechanism sometimes gives no response, there will be distinct properties p and q such that both are present when it responds positively, both are absent when it responds negatively, but sometimes one is present and the other absent when it does not respond, and that since it is equally good at registering p and q , and no better at registering any other property, the question "Which property does this mechanism best register?" has no unique answer. This objection ignores the statistical nature of reliability. The mechanism cannot be expected to register any distal property infallibly; since its functioning depends on the state of the subject as well as on the state of the environment, no distal property will be present whenever there is a positive response and absent whenever there is a negative one.¹⁶

15. A more teleological question would be "Which property did this mechanism evolve to register?". Considerations like those in the text would still apply.

16. Why consider distal properties rather than proximal ones? This is a general but not unanswerable question for causal theories of reference; it is not a special problem for the epistemic view of vagueness.

Reliability is a matter of minimizing a non-zero probability of error; for all that has been shown, just one property may do that.¹⁷

A subject whose primary access to a property is through a recognitional mechanism may not be helped to detect it by extra information of a kind which cannot be processed by that mechanism, even if the new information is in fact a reliable indicator of the presence of the property — for the subject may not know that. My exact measurements may in fact be a sufficient condition for thinness, and knowledge of the former still not enable us to derive knowledge of the latter; for all that, thinness may be the property best registered by our perceptual recognitional capacity for thinness.

The foregoing speculations should not mislead one into supposing that a causal theory of reference is essential to an epistemic view of vagueness. They illustrate only one way in which our use of a vague expression might determine a sharp property. A comprehensive account of the connection between meaning and use would no doubt be very different. Since no one knows what such an account would be like, the epistemic view of vagueness should not be singled out for its failure to provide one. No reason has emerged to think that it makes such an account harder to provide. At the worst, there may be no account to be had, beyond a few vague salutatory remarks. Meaning may supervene on use in an unsurveyably chaotic way.

5 The charge against the epistemic view of vagueness might be revised. If the view does not force what we mean to transcend what we do, perhaps it forces what we mean to transcend what we know. The new charge is as obscure as the old one, but may be worth exploring.

A cautious answer is that the epistemic view of vagueness allows us to know what we mean. No gap need open between what we mean and what we think we mean, for both are determined in the same way, perhaps that described in section 4. We know that "TW is tall" as we use it means that, and is true if and only if, TW is tall. If we cannot know whether TW is tall, who but the verificationist thought that actual knowledge of truth conditions requires possible knowledge of truth value?

It may be replied that the epistemic view makes us ignorant of the sense of a vague term, not just of its reference. Of course we do not know where all the thin things are in physical space; the point is that we should not even know where they all are in conceptual space. We should be using a term that does in fact determine a line in conceptual

17. If several properties tie for first place, the obvious candidate is their conjunction (even if it is not itself one of them).

space without being able to locate that line. We should understand it partially, as one partially understands a word one has heard used once or twice. But in the latter case the word's meaning is backed by other speakers' full understanding, whereas no one is allowed full understanding of the vague term. The objection to the epistemic view is that it attributes partial understanding to the speech community as a whole. It is not entitled to say that we know what we mean. It attributes to the community incomplete knowledge of a complete meaning; would it not be more reasonable to attribute complete knowledge of an incomplete meaning?

The objection is based on the Fregean model of the sense of a term as a region in conceptual space: to grasp a sense is to know where its boundary runs. Individual points in this space are located by means of precise descriptions such as "having exact physical measurements m ". Thus the demand that one know which points are in the region marked off by a vague term such as "thin" is simply the demand that one know truths such as "Anyone having exact physical measurements m is thin" or "No one having exact physical measurements m is thin". The unreasonableness of that demand was already noted in section 3; the metaphysical necessity of such truths does not justify the demand to know them. The metaphor of conceptual space adds no force to the demand. Rather, its function is illicitly to collapse distinctions between concepts whose equivalence is metaphysically necessary but not *a priori*. If a proposition is identified with a region in a space of possible worlds, cognitively significant distinctions are lost in a familiar way; exactly the same happens when the objection identifies a sense with a region in conceptual space.

On the epistemic view, our understanding of vague terms is not partial. The measure of full understanding is not possession of a complete set of metaphysically necessary truths but complete induction into a practice. When I have heard a word used only once or twice, my understanding is partial because there is more to the community's use of it than I yet know. I have not got fully inside the practice; I am to some extent still an outsider. It does not follow that if we had all understood the term in the vague way I do, all our understandings would have been partial, though they would still have determined complete intensions.¹⁸ In that counterfactual situation, we should all have been insiders. To know what a word means is to be completely inducted into a practice that does in fact determine a complete intension.

18. My deference to speakers with fuller understanding may be excluded from the counterfactual situation. Think of an intension as a function from possible worlds to extensions.

That rather minimalist answer to the objection is enough. However, a more speculative line of thought may be mentioned. If meaning supervenes on use, might it also supervene on knowledge? The idea can be developed. Let the *verification conditions* of a sentence be those in which its truth conditions knowably obtain, and its *falsification conditions* be those in which its truth conditions knowably fail to obtain. A kind of supervenience claim quite consistent with the epistemic view of vagueness is:

(@) If two sentences have the same verification conditions and the same falsification conditions, they have the same truth conditions.

(@) claims a supervenience of truth conditions on verification and falsification conditions. It no more identifies truth conditions with verification conditions than it identifies them with falsification conditions. In general, (@) is probably too strong. For example, there may be a sentence whose truth conditions cannot be known to obtain and cannot be known not to obtain; it would have the same verification conditions and falsification conditions as its negation, but not the same truth conditions. However, ordinary vague sentences are not like that. (@) might hold for them. In fact a formal version of (@) can be proved for a simple modal logic in which "necessity" is interpreted as knowability, truth does not entail knowability, and the underlying propositional logic is classical.¹⁹

(@) will not satisfy reductionist aspirations, for the truth conditions are used in characterizing the verification and falsification conditions. But that is a problem for the reductionist aspirations, not for the epistemic view of vagueness. What the consistency of (@) with the epistemic view shows is that the latter does not force what we mean to transcend what we know, if the purport of the charge is that the epistemic view would not allow truth conditions to supervene on the conditions in which they can be known to obtain or not to obtain.

6 Little has been said to explain our ignorance in borderline cases. Of course, ignorance might be taken as the normal state: perhaps we should think of knowledge as impossible unless special circumstances make it possible, rather than as possible unless special circumstances make it impossible. However, we may be able to do better than that in the case at hand.

Consider again the supervenience of meaning on use, at least for a fixed contribution from the environment. For any difference in meaning,

19. See my 1990b. The result is consistent with the doubt expressed about (@), since it cannot automatically be lifted to extensions of the language.

there is a difference in use. The converse does not always hold. The meaning of a word may be stabilized by natural divisions, so that a small difference in use would make no difference in meaning. A slightly increased propensity to call fool's gold "gold" would not change the meaning of the word "gold". But the meaning of a vague word is not stabilized by natural divisions in this way. A slight shift in our dispositions to call things "thin" would slightly shift the meaning of "thin". On the epistemic view, the boundary of "thin" is sharp but unstable. Suppose that I am on the "thin" side of the boundary, but only just. If our use of "thin" had been very slightly different, as it easily could have been, I would have been on the "not thin" side. The sentence "TW is thin" is true, but could easily have been false.²⁰ Moreover, someone who utters it assertively could easily have done so falsely, for the decision to utter it was not sensitive to all the slight shifts in the use of "thin" that would make the utterance false.

The point is not confined to public language. Even idiolects are vague. You may have no settled disposition to assent to or dissent from "TW is thin". If you were forced to go one way or the other, which way you went would depend on your circumstances and mood. If you assented, that would not automatically make the sentence true in your idiolect; if you dissented, that would not automatically make it false. What you mean by "thin" does not change with every change in your circumstances and mood. The extension of a term in your idiolect depends on the whole pattern of your use in a variety of circumstances and moods; you have no way of making each part of your use perfectly sensitive to the whole, for you have no way of surveying the whole. To imagine away this sprawling quality of your use is to imagine away its vagueness.²¹

An utterance of "TW is thin" is not the outcome of a disposition to be reliably right; it is right by luck. It can therefore hardly be an expression of knowledge. Contrastively, an utterance of "TW is thin" is an expression of knowledge only if I am some way from the boundary of "thin", that is, only if anyone with physical measurements very close to mine is also thin. More generally, for a given way of measuring difference in physical measurements there will be a small but non-zero constant c such that:

- (!) If x and y differ in physical measurements by less than c and x is known to be thin, y is thin.

20. The point is not that I might easily not have been thin. In the relevant counterfactual situations, my physical measurements are just what they actually are, but "thin" means something slightly different from what it actually means.

21. What goes for words in your idiolect also goes for your concepts.

Similar principles can be formulated for other vague terms. Vague knowledge requires a margin for error.

Given (!), one cannot know a conjunction of the form " x is thin and y is not thin" when x and y differ in physical measurements by less than c . To know the conjunction, one would have to know its first conjunct; but then by (!) its second conjunct would be false, making the whole conjunction false and therefore unknown. Since such conjunctions cannot be known, the unwary may suppose that they cannot be true. "Thin" will then look as though it is governed by a tolerance principle of the form: if x and y differ in physical measurements by less than c and x is thin, y is thin. One can now construct a sorites paradox by considering a series of men, the first very thin, the last very fat, and each differing from the next in physical measurements by less than c : by repeated applications of the tolerance principle, since the first man is thin, so is the last man. Fortunately, "thin" is not governed by the tolerance principle; it is governed by the margin for error principle (!), which generates no sorites paradox.²²

The plausibility of (!) does not depend on the epistemic view of vagueness. Its rationale is that reliable truth is a necessary (perhaps not sufficient) condition of knowledge, and that a vague judgement is reliably true only if it is true in sufficiently similar cases. This point does not require the judgement to be true or false in every case. But once our uncertainty has been explained in terms of the independently plausible principle (!), it no longer provides a reason for not asserting bivalence, for bivalence is quite compatible with (!).

7 The most obvious argument for the epistemic view of vagueness has so far not been mentioned. The epistemic view involves no revision of classical logic and semantics; its rivals do involve such revisions. Classical logic and semantics are vastly superior to the alternatives in simplicity, power, past success, and integration with theories in other domains. In these circumstances it would be sensible to adopt the epistemic view in order to retain classical logic and semantics even if it were subject to philosophical criticisms in which we could locate no

22. (!) might be thought to generate a sorites paradox not for "thin" but for "known to be thin", given that (!) is known, that each man in the series is known to differ from the next in physical measurements by less than c , and that the very thin man is known to be thin. However, analysis of the argument shows it to require the KK principle that what is known is known to be known. But since "known to be thin" is itself vague, it too obeys a margin for error principle, which in turn implies that one can know x to be thin without being in a position to know that one knows that x is thin. Thus the KK principle fails. The failure of the KK principle (i.e. the S4 axiom) in the modal logic KT is essential to the result cited at n. 19.

fallacy; not every anomaly falsifies a theory.²³ Although that second line of defence exists, there is no need to occupy it if the argument of this paper is correct, for we can locate the fallacies in philosophical criticisms of the epistemic view of vagueness.²⁴

15 Sorites paradoxes and the semantics of vagueness

Michael Tye

It is sometimes supposed that sorites paradoxes are an inevitable consequence of the very nature of vagueness. Take, for example, the term "bald". If "bald" is vague then it lacks precise boundaries. So

- (1) There is a definite number, N , such that a man with N hairs on his head is bald and a man with $N + 1$ hairs on his head is not

is false. But intuitively the denial of (1) is equivalent to the assertion of

- (2) For any definite number, N , if a man with N hairs on his head is bald then a man with $N + 1$ hairs on his head is also bald.

And (2), together with the obvious truth

- (3) A man with no hairs on his head is bald

entails the obvious falsehood

- (4) A man with a million hairs on his head is bald

via a million applications of modus ponens and universal instantiation. To treat this line of reasoning as a *reductio* of the denial of (1) is to concede that "bald" is not vague, and hence, in the general case, to concede that no predicates are vague. This conclusion is, of course, itself paradoxical. What, then, has gone wrong?

In this paper I want to present a novel semantics of vagueness which is, I maintain, invulnerable to sorites paradoxes such as the one above. On the approach I favor, there are three truth-values: true, false, and neither true nor false (or indefinite). The third value here is, strictly speaking, not a truth-value at all but rather a truth-value gap. In my view, there are gaps due to failure of reference or presupposition and gaps due to vagueness.¹ Corresponding to the two-valued connectives

From *Philosophical Perspectives* 8: *Logic and Language*, ed. James E. Tomberlin (1994), pp. 189–206. Atascadero, CA: Ridgeview Publishing Co. This selection omits §1 of the original paper. © Ridgeview Publishing Co. Reprinted by permission.

1. Where a gap is due to vagueness, I maintain that something is said which is neither true nor false. I deny, however, that anything is said in the case where a gap is due to failure of reference. I am inclined to extend the latter view to gaps due to failure of presupposition.

23. For another argument for the epistemic view see my 1990a, p. 107.

24. Some of the material on which this paper is based has been presented in talks at the universities of Oxford, London (University College), Dundee, Stirling, A.N.U., New England (Armidale), Queensland, Monash, Bradford, Lisbon, and Bristol. More people have helped with good questions than I can name.