

THE EPISTEMIC THEORY OF VAGUENESS

Stephen Schiffer
New York University

I. Introduction

Suppose that Harry is a borderline case of baldness. Then *the epistemic theory of vagueness* has it that it's either true that he's bald or else true that he's not bald, but nothing we do will ever enable us to know the truth about Harry's baldness; and likewise, *mutatis mutandis*, for every other borderline case of a vague notion. This remarkable thesis is defended with great force and ingenuity by Timothy Williamson in his masterful book *Vagueness*,¹ but several other extremely able contemporary philosophers also accept the theory in the sense in which I'm about to define it, and they include Roy Sorensen, Paul Horwich, Hartry Field, Vann McGee and Brian McLaughlin.² Consequently, this paper will focus not only on Williamson's version of the epistemic theory, but also on the theory in its other guises. More specifically, this paper has the following outline.

- A definition of the epistemic theory in the sense in which I want to discuss it.
- A brief discussion of the motivation for the generic epistemic theory.
- Application of the epistemic theory to the two issues that define the philosophical problem of vagueness—the problem of resolving the sorites paradox and the problem of explicating the notion of a borderline case (and, thereby, the notion of vagueness, for vagueness just is the possibility of borderline cases).
- The outstanding question for the epistemic theorist is how to explain the ignorance to which she's committed. First I'll discuss how this challenge might be met by those epistemic theorists, such as Williamson, who take the crucial semantic properties to be use dependent. Then I'll discuss how the challenge might be met by those epistemic theorists, such as Hartry Field, who take the crucial semantic properties to be use independent.

- A concluding suggestion about a competing approach to the problems of vagueness.

II. The Epistemic Theory Defined and Motivated

Timothy Williamson, in a précis of *Vagueness*, defines the epistemic theory thus: it's the thesis

that the proposition a vague sentence expresses in a borderline case is true or false, and we cannot know which. We are ignorant of its truth-value.³

Alternatively, but to the same effect, we may define the epistemic theory as the thesis that:

- (1) there are vague propositions—the actual or potential propositional contents of utterances of vague sentences;
- (2) bivalence holds for all vague propositions (i.e., every vague proposition is either true or false);⁴
- (3) true, but borderline, vague propositions are unknowable (i.e., we can't know the truth-values of vague propositions when they describe borderline cases).

This can stand some comment.

The notion of a proposition invoked in (1) is entirely neutral as regards ontological commitment to things of any kind, let alone to mind- and language-independent abstract entities. I'm simply alluding to the pleonastic equivalence between

A believes that *S*

and

A believes the proposition that *S*.

In other words, 'the proposition' in 'the proposition that *S*' is semantically superfluous. But it's important to define the epistemic theory in terms of whatever is going on with that-clauses, because the epistemic theory is a theory of vagueness as *ignorance*, and both knowledge and failure of knowledge always concern one's knowing, or failing to know, *that such-and-such*. What we can't know, if the epistemic theory is right, is *that Harry is/isn't bald*.

Yet, notwithstanding his above-quoted definition of the epistemic theory, Williamson, in *Vagueness*, argues that the principle of bivalence needed by the epistemic theory is *not* bivalence for propositions. He says:

[the principle that every proposition is either true or false] does not bear very directly on problems of vagueness. A philosopher might endorse bivalence for propositions, while treating vagueness as the failure of an utterance to express a unique proposition. On this view, a vague utterance in a borderline case expresses some true propositions and some false ones (a form of supervaluationism might result). There is no commitment to a bivalent classification of utterances, or to the ignorance on our part that such a classification implies. The problem of vagueness is a problem about the classification of utterances. To debate a form of bivalence in which the truth-bearers are propositions is to miss the point of controversy.⁵

He then proposes that the required principle of bivalence “claims truth or falsity when, and only when, something has been said to be the case,”⁶ and formulates the principle of bivalence as the following schema:

[B] If *u* says that *P*, then either *u* is true or *u* is false.⁷

[B] in turn is classically entailed by Tarski-like schemas Williamson takes to be definitive of utterance truth and falsity:

If *u* says that *P*, then *u* is true if and only if *P*.
If *u* says that *P*, then *u* is false if and only if not *P*.⁸

It’s true that the principle of bivalence for propositions won’t itself yield the epistemic theory we want—that’s already conceded by the above definition of the theory whose only bivalence principle governs propositions. But not only *can* the epistemic theory be formulated in terms of bivalence for propositions, it *must* be so formulated. For, as I’ve already said, the epistemic theory entails that one is *ignorant* of the truth-value of *what is said* by a vague utterance in a borderline case, and there can be no such *ignorance* unless the propositions expressed by such utterances have truth-values of which one is ignorant. The problem with Williamson’s principle [B] is that one can accept it while also denying that *what is said* by a vague utterance in a borderline case is true or false.

To see this it’s enough to notice that there’s a consistent position which denies that bivalence holds for the *propositions* asserted by vague utterances in borderline cases but yet, at the same time, accepts both classical logic and Williamson’s [B]. The theorist of this position might hold the following two things:

- (a) A supervaluationist account of *propositions*, whereby excluded middle holds but where neither disjunct in a proposition of the form *P or not-P* is true when those disjuncts are borderline propositions. In other words, if Harry is a borderline case of baldness, then the proposition *that Harry is bald or Harry isn’t bald* is true, although neither the proposition *that Harry is bald* nor the proposition *that Harry isn’t bald* is true. This theorist, needless to say, rejects the following two schemas for propositions:

The proposition that S is true iff S .

The proposition that S is false iff not S .

- (b) A deflationary account of utterance truth or falsity whereby it's held that 'true' and 'false' as applied to utterances are purely disquotational predicates, such that what is said by an utterance of ' u is true' is the same as what is said by u , and what is said by an utterance of ' u is false' is the same as what is said by an utterance of the negation of the sentence produced in u .⁹

A theorist holding (a) and (b) will be committed to accepting Williamson's disquotational principles of truth and falsity (for those principles, for her, reduce to the tautologies *If u says that P , then P iff P* and *If u says that P , then not P iff not P*), but she will deny that the truth-value of an utterance is determined by that of the proposition it expresses. It's clear that this imagined theorist won't be an epistemic theorist. The epistemic theory holds that it's either true or false *that borderline Harry is bald*, which is precisely what's denied by the position in question. According to this position, one can't know that Harry is bald and one can't know that he's not bald, but this doesn't constitute *ignorance*, because it's neither true that Harry is bald nor true that he's not bald.¹⁰

One reason I have for pressing what's apt to seem a nit-picking point is to clarify the status of a certain observation that is often made in supervaluationist discussions of vagueness. The observation is that anyone—supervaluationist of course included—who accepts classical logic can trivially have bivalence for sentences or utterances by helping himself to a disquotational sense of 'true' and 'false'. Thus, Kit Fine writes:

I think that [the definitely-operator] is a prior notion to 'true' and not conversely. For let ' true_T ' be that notion of truth that satisfies the Tarski-equivalence, even for vague sentences:

' A ' is true_T if and only if A .

The vagueness of ' true_T ' waxes and wanes, as it were, with the vagueness of the given sentence; so that if a denotes a borderline case of F then Fa is a borderline case of ' true_T '. Then the ordinary notion of truth is given by the definition:

x is true =_{df} Definitely (x is true_T).

Thus ' true_T ' is primary; 'true' is secondary and to be defined with the help of the definitely-operator.¹¹

And Vann McGee and Brian McLaughlin write:

The practical advantages of having a notion of truth that satisfies the (T)-sentences have been stressed by Quine and others... . Happily, these advantages will still be

available to us, even if we adopt a [non-bivalent] conception. While we allow our usage of the word ‘true’ to be governed by [e.g. a supervaluationist conception of truth as super-truth], we can introduce a new predicate ‘plue’ (an elision of ‘pleonastically true’), implicitly defined by the axiom schema:

(P) *P* is plue if and only if P.¹²

The simple point I’m concerned to stress is that the bivalence crucial to the epistemic theory can’t trivially be had by any theorist who accepts excluded middle. If one is to be a vagueness-*as-ignorance* theorist, one must hold that vague utterances in borderline cases express *propositions*—that is to say, say something—for which bivalence holds.

It may be thought that even the definition of the epistemic theory which uses bivalence for propositions isn’t really enough to make one *truly* an epistemic theorist. For can’t one always introduce predicates of propositions, ‘true_p’ and ‘false_p’, such that the following two schemas are analytically valid?

The proposition that *S* is true_p iff *S*.

The proposition that *S* is false_p iff not *S*.

And if so, then one has bivalence_p if one has excluded middle.

Now, of course one can introduce such predicates via stipulations about the displayed schemas, but here we need to ask how these newly defined properties of propositions relate to the properties of truth and falsity of propositions. If they’re different, nothing has been gained or lost. Knowledge, and hence ignorance, is analytically tied to the *truth* of propositions, whatever that property comes to. As theorists, we’re not starting from scratch, so that we’re free to ask what notion of truth for propositions best suits us. Rather, we begin with the notion of a proposition and the notion of a proposition’s being true, and it’s in terms of those notions that our concepts of knowledge and meaning must be understood. This is especially so since the notions of propositional truth and falsity we employ are apt to be understood specifically in terms of features of the propositions to which those notions apply. For example, if propositions are ordered pairs of the form $\langle\langle x_1, \dots, x_n \rangle, \Phi^n\rangle$, where $\langle x_1, \dots, x_n \rangle$ is an *n*-ary sequence of items and Φ^n is an *n*-ary property, then we’ll want to say that such a proposition is true just in case its *n*-ary sequence instantiates its *n*-ary property and false just in case it fails to instantiate it.

It’s not contentious to define the epistemic theory in terms of propositions in the sense of “whatever is ascribed in that-clauses,” but clause (1) of the definition nevertheless makes a potentially contentious claim. This is the claim that there are *vague* propositions. For this claim is tantamount to the claim that an utterance of ‘Harry is bald’ determinately expresses the vague proposition that Harry is bald,¹³ and a theorist, as Williamson himself remarked in an above-cited passage, might want to claim that there are no vague propositions but only absolutely

precise ones and that what makes an *utterance* vague is that it indeterminately, or partially,¹⁴ expresses each of an array of precise propositions. This would then doubtless be joined to a supervaluationist semantics whereby an utterance of a vague sentence is true just in case each of the precise propositions it partially expresses is true, false just in case each of the precise propositions is false, and neither true nor false just in case some of them are true while others are false.

Fortunately for the epistemic theorist, however, there is a compelling argument against the precise-proposition line and, thereby, in favor of the vague-proposition line her theory requires. The argument, run on an arbitrary example, goes as follows. If Louise, speaking sincerely and literally, utters ‘Harry is bald’, then it’s true—determinately true, if you like—that Louise said that Harry is bald. Now, for all intents and purposes, there are as regards the that-clause in

[*] Louise said that Harry is bald

only two options: either it determinately refers to the vague proposition that Harry is bald or it partially, or indeterminately, refers to each of an array of precise propositions. But we can rule out the second option thus: (i) If the that-clause partially refers to each of an array of precise propositions, then [*] is true just in case it comes out true no matter which of the precise propositions is taken to be the reference of its that-clause. (ii) But then [*], which is true, would fail to be true, for, even allowing for the vagueness of ‘said that’, Louise didn’t say *any* of the precise propositions in question (e.g., she didn’t say that Harry has fewer than 3,832 hairs on his scalp). (iii) Therefore, [*]’s that-clause doesn’t partially refer to each of an array of precise propositions and must instead refer to the vague proposition that Harry is bald.¹⁵

So it’s very reasonable to suppose that clause (1) of the definition of the epistemic theory [there are vague propositions] is true. Clause (3) of the definition [true borderline propositions are unknowable] is obviously true, if clause (2) [vague propositions are bivalent] is true: if the proposition that borderline Harry is bald has a truth-value, it’s obvious we can’t know what it is. It follows that the epistemic theory enjoys very strong support if it’s plausible that clause (2) of its definition is true. We need, then, to ask what reasons there are for thinking that bivalence holds for vague propositions. Since it’s reasonable to suppose that bivalence holds for vague propositions only if it’s reasonable to suppose it holds for all propositions, we need to ask what reasons there are for supposing bivalence to hold for all propositions.¹⁶

Well, (2) enjoys all the support that can be mustered for the plausible conjunction that the law of excluded middle (every instance of the schema *the proposition that p or not-p is true* is true) is true and every instance of the following schema is true:

The proposition that *S* is true (more colloquially: it’s true that *S*) iff *S*.

Nevertheless, unless there is something new to be said for the conjunction, it's not polemically powerful against those who would deal with vagueness by not accepting it.

III. The Epistemic Theory and the Two Issues of Vagueness

If the epistemic theory is correct, then it ought to give a satisfactory resolution of the two issues that define the philosophical problem of vagueness, viz. the problem of explicating the notion of a borderline case and the sorites paradox.

As already noticed, the problem of saying in what being a borderline case consists is actually the problem of saying in what vagueness consists, since vagueness just is the possibility of borderline cases. Vagueness can infect any kind of term, but to simplify the present discussion I'll limit myself to predicate vagueness.

If the epistemic theory were correct, it would provide a definition of being a borderline case. An epistemic theorist of any stripe is committed to the truth of some instance of this definition scheme:

x is a borderline case of *F* iff it's either true that *x* is *F* or true that *x* isn't *F*, but it's impossible for anyone to know, *for such-and-such reason*, the truth about *x*'s being *F*.

Without the italicized qualification, the right-hand side would fail to provide a sufficient condition for being a borderline case. It may be that no one can know the truth-value of Goldbach's conjecture, but nothing about vagueness follows from that. For ignorance to imply borderlineness, it must have a particular source. Now, each epistemic theorist is committed to explaining the ignorance to which she's committed, and whatever that explanation is, it will enable the theorist to say what sort of ignorance is definitive of vagueness. Therefore, the epistemic theorist can resolve the first problem of vagueness, *if* she can satisfactorily explain the ignorance to which she's committed.

So much for the epistemic theory's ability to explicate the notion of a borderline case. Let's turn to what it can do to resolve the sorites paradox. The sorites can take several interestingly different forms, but in order to simplify the discussion, I'll consider only the paradox that arises when we consider the following set of propositions:

- (1) A person with \$50,000,000 is rich.
- (2) For any *n*, if a person with \$*n* is rich, then so is a person with \$*n* - \$1.
- (3) It's not the case that a person with only \$3 is rich.

Obviously, this presents a paradox because, on the one hand, the three propositions are mutually inconsistent by classical logic, as (1) and (2) classically entail the negation of (3), and, on the other hand, each of the propositions appears plausible when considered on its own. Assuming that mutually inconsistent prop-

ositions can't all be true, one may reasonably suppose that a fully adequate solution to the paradox would do two things: first, tell us which of (1)-(3) isn't true; and second, explain to us why it appears to be true, so that it will be stripped of its patina of truth, never again to tempt us. As regards the first task, the epistemic theory tells us that (2) is the odd-guy-out. As regards the second task, it tells us that (2) is not merely not true, but false, and its falsity entails that there is some specific sharp \$1 cutoff between the rich and the non-rich. That is to say, there is some numeral α such that *Having $\$ \alpha$ is sufficient for being rich, but having $\$ \alpha - \1 isn't* is true. No doubt the theorist will tell us that at least part of the reason we think there is no such sharp cutoff is that it's impossible for us to know what it is. Whether this is part of a good solution remains to be considered.

IV. The Epistemic Theory's Big Debt

The epistemic theory sorely needs an explanation of the ignorance it entails. Before looking at possible explanations, we need to get clear as to exactly what sort of ignorance is entailed. Obviously, the theory entails that we can't know the truth-value of the proposition expressed by a vague utterance in a borderline case. If you say that Harry is bald when Harry is a borderline case of baldness, then what you said is either true or else false, but no one can know the truth-value of your statement. Of course, the epistemic theorist is committed to ignorance about borderline cases even when no vague statements are produced. If Harry's a borderline case of baldness, then it's either a fact that Harry is bald or else a fact that he's not bald, and no one can know what the fact of the matter about Harry's baldness is.

Actually, there will be many slightly different baldness facts about borderline Harry no one can hope to know. This is because the epistemic theory must allow that while utterances of 'bald' express determinate properties, the properties expressed can vary somewhat from context to context. This concession is needed to accommodate the fact that the penumbras of vague terms can dilate or constrict depending on the circumstances of utterance; indeed, the epistemic theorist must allow that an utterance of 'Harry is bald' may be true in one context and false in another. This context sensitivity can for the most part be ignored if we suppose, with Williamson, that each occurrence of a vague sentence expresses a determinate proposition.

The ignorance the epistemic theorist must explain extends to more than the application of vague properties. Suppose that in describing a certain incident Al says 'Betty was standing roughly there'. Provided Al said something, what he said was either true or false, whether or not this vague utterance describes a borderline case. This entails that there is some absolutely precise region of space such that the utterance is true just in case Betty was standing within that region, false otherwise. When the utterance describes a borderline case, no one can know the truth-value of what was said. Likewise, the epistemic theorist must hold that when I utter 'I worked for a little while yesterday' and state thereby that I worked

for a little while yesterday, my utterance is true just in case the time I spent working falls within a certain absolutely precise span of time and is false just in case the amount of time I spent working was even one nanosecond longer. When the utterance describes a borderline case, no one can know the truth-value of what it says. And when you uttered 'I put a pinch of salt in the sauce', your utterance of 'a pinch of salt' refers to a precise range of quantities of salt, such that your statement is false just in case the amount of salt you put in was so much as one-millionth of a milligram outside that range. When the utterance describes a borderline case, no one can know the truth-value of what it says. The epistemic theorist must explain why it's impossible for us to know the truth-values of the precise propositions said in these vague borderline utterances.

There are metaphysically necessary general truths to whose unknowability the epistemic theory is committed. They have to do with the supervenience of vague properties on other properties. For consider baldness. This property supervenes on the hair situation on a person's scalp—on the number, size, distribution, etc., of the hairs on a person's scalp—in that, necessarily, if a person is bald (not bald), then so is anyone with exactly the same hair situation on his or her scalp. For convenience of exposition, let's harmlessly simplify and pretend that baldness supervenes just on the number of hairs on a person's scalp, and let's ignore the vagueness of 'hair' and 'scalp'. Relative to this simplification, it will follow that there is some number n such that it's a metaphysically necessary truth that a person is bald iff he or she has fewer than n hairs on his or her scalp. Let's take that number to be 3,832. Then, since having fewer than 3,832 hairs on one's scalp is a metaphysically necessary and sufficient condition for being bald, we can for all intents and purposes say that baldness = the property of having fewer than 3,832 hairs on one's scalp. But the epistemic theory must hold that no one can know this metaphysically necessary identity proposition. For if one could, one could count the hairs on borderline Harry's scalp and thereby come to know whether or not he's bald. So this is another sort of ignorance the epistemic theory must explain.

Finally, there are semantic facts to whose unknowability the epistemic theory is committed. Suppose Al says of borderline Harry 'There's no saying whether or not he's bald'. Al's utterance of 'bald' refers to baldness, and this is something we and Al know. Since baldness = the property of having fewer than 3,832 hairs on one's scalp, Al's utterance of 'bald' refers to the property of having fewer than 3,832 hairs on one's scalp. But the epistemic theorist must hold that no one can know that fact. For if one could, one could infer that baldness = the property of having 3,832 hairs on one's scalp, and one could then resolve borderline cases of baldness by counting hairs. So this is another kind of ignorance the epistemic theorist must explain. Likewise as regards what the epistemic theorist must say about the vague singular reference in Al's utterance of 'Betty was standing roughly there' and about the reference of 'a pinch of salt' in your utterance of 'I put a pinch of salt in the sauce'. The theorist must hold that the utterance of 'roughly there' refers to an absolutely precisely bounded region of space and that the utterance of

'a pinch of salt' refers to an absolutely precise range of quantities of salt, but whether or not one can know the truth-values of the propositions expressed in these utterances, no one can know the exact boundaries of the region of space or the precise range of quantities of salt to which references were made. And what applies to the reference of public language utterances applies as well, *mutatis mutandis*, to the references of our mental states, of, if you will, the expressions in our inner system of mental representation.

V. Explaining the Ignorance: Use-Dependent Semantic Properties

Epistemic theorists may be divided into those for whom crucial semantic properties are use-*dependent* properties and those for whom they're use-*independent* properties. A semantic property is use dependent provided that whether or not an expression instantiates it depends on how that expression is *used*. For example, it's reasonable to hold that the name-of relation is a use-dependent semantic property, since it's reasonable to suppose that 'Timothy Williamson' names Timothy Williamson by virtue of the way the name 'Timothy Williamson' is used by certain people. A semantic property is use independent provided that it's not use dependent. Those who wear the epistemic label proudly, such as Roy Sorensen and Timothy Williamson, hold that semantic properties are use dependent, whereas those, such as Hartry Field, Vann McGee and Brian McLaughlin, who wear the label reluctantly, hold that at least certain crucial semantic properties are use independent. Let's start with those on the use-dependent side.

Roy Sorensen holds that no explanation is needed.¹⁷ *Knowledge*, he claims, always needs an explanation, but not the *absence* of knowledge. I don't find this plausible. It seems not to be based on any plausible principle of knowledge, and we often demand and expect an explanation for a person's failure to know: How could you not know that you're scheduled to give a talk today at the University of Transylvania? I won't discuss this option further.

Timothy Williamson is the epistemic theorist on the use-dependent side with the most fully elaborated explanation. According to him, all the ignorance to which the epistemic theory is committed is to be explained in the same way, as ignorance that "is just what independently justified epistemic principles would lead one to expect."¹⁸ The allusion is to what Williamson calls *margin-for-error principles*, which he claims govern cases of inexact knowledge. But it's possible to give Williamson's account of the ignorance to which he's committed without explicitly mentioning inexact knowledge or the principles that govern it. The crucial idea is as follows. For a true belief to count as knowledge, the mechanism that produced it must be reliable. For the propositions in question, there are no belief-forming mechanisms reliable enough to yield knowledge. Whatever belief-forming mechanism might produce a belief of the relevant kind which happens to be true, it would have produced that belief even if the belief had had a very slightly different content but one that made it false. A mechanism that might

result in false beliefs as easily as in true beliefs isn't reliable enough to yield knowledge.

Let me illustrate how this explanation works with respect to some of our ongoing examples. For vividness, I'll pretend that our subject, Jane, thinks in neural English. Jane believes a proposition p when, and only when, a sentence of her *lingua mentis* which means p is tokened in her "belief box." Harry continues to be our paradigm borderline case of baldness, but we'll suppose that he is in fact bald, since baldness is the property of having fewer than 3,832 hairs on one's scalp and he happens to have exactly 3,831 on his scalp. 'Bald' in Jane's neural English refers to baldness, i.e., to the property of having fewer than 3,832 hairs on one's scalp. The reference relation that relates 'bald' in Jane's head to baldness is determined by various factors—causal/environmental factors, the conceptual roles of the terms in its domain, and perhaps also by the way the public language counterparts of those terms are used in a thinker's linguistic community. This relation is super sensitive to variations in some of those factors; changes in those factors that are indiscernible to normal thinkers can make 'bald' in Jane's head change its reference from baldness, the property of having fewer than 3,832 hairs on one's scalp, to baldness', the property of having fewer than 3,831 hairs on one's scalp, or to any of the other properties casual inspection can't distinguish from baldness. This puts us in a position to explain why Jane can't possibly know that Harry is bald, even though he is. The explanation is simply that whatever belief-forming mechanism might put a token of 'Harry is bald' in Jane's belief box, thereby giving her the true belief that Harry is bald, it would have put that sentence there even if 'bald' referred not to baldness, but to baldness', the property of having fewer than 3,831 hairs on one's scalp. Thus, the mechanism might just as well have produced in Jane the *false* belief that Harry is bald', and therefore isn't reliable enough to yield knowledge. The same sort of explanation explains why Jane can't know the metaphysically necessary truth that a person is bald just in case he or she has fewer than 3,832 hairs on his or her scalp.

But what explains the fact that Jane can't know that 'bald' refers to the property of having fewer than 3,832 hairs on one's scalp? The proposition that the word 'bald' refers to the property of having fewer than 3,832 hairs on one's scalp doesn't involve the concept of baldness, so this ignorance can't be explained by the vagueness of 'bald'. To this Williamson will respond that the ignorance, though not explained by the vagueness of 'bald' in Jane's *lingua mentis*, is explained by the vagueness of the verb 'to refer' in her *lingua mentis*. The explanation from there takes the course already sketched.¹⁹

I have at least three problems with Williamson's explanation of the ignorance to which he's committed. The first is really a problem with any version of the epistemic theory, but I'll mention it here to get it out of the way. The epistemic theory hopes to provide what I've elsewhere called a *happy-face solution* to the sorites.²⁰ Such a solution, as already indicated, hopes not merely to tell us which of our three mutually incompatible sorites propositions (see above p. 487) isn't true, but also hopes to explain why it looked to us to be true. At least part of that

explanation, presumably, is that, although there is a precise \$1 cutoff, we can't possibly know where it falls. This is at best a partial explanation because a further explanation is needed as to why we should take failure to know to mean that there's nothing there to know. For all we know, it's impossible to prove or disprove Goldbach's conjecture that every even number greater than or equal to 6 is the sum of two odd primes, but most people who were convinced of such an impossibility would conclude that there was here a fact that couldn't be known. Why, then, should we conclude that there is no \$1 cutoff separating the rich from the non-rich just because we can't find such a cutoff? There must be more to the epistemic theorist's explanation of why we're mistakenly tempted by the sorites premise, but, to the best of my knowledge, no epistemic theorist has even offered a completion. In fact, our inability to find a cutoff in and of itself *isn't* why we find it incredible that there is some definite number—say, 639,472—such that having that many dollars suffices to make a person rich while having one dollar less than that doesn't. Rather, I submit, our wanting to deny there's a \$1 cutoff is because it strikes us as incredible that our use of 'rich' should determine such a cutoff. This observation brings us to the second problem.

The second problem has to do with the fact that Williamson's account of ignorance entails that each vague expression is so used that that use, together with the environment in which the expression is used, determines an absolutely precise reference for that expression. As Vann McGee and Brian McLaughlin put it,

According to [Williamson's elaboration of] the epistemic theory, the thoughts and practices of English speakers, together with an exact specification of the number and configuration of hairs on a person's head—a specification that could be given in terms of Cartesian coordinates—will suffice to determine whether that person satisfies "bald" or "not bald." The thoughts and practices of English speakers determine an exact number r that is the critical point separating those parcels that satisfy "weighs a little less than ten pounds" from those that satisfy its negation.²¹

This is problematic for two reasons. First, Williamson has no direct reason stemming from views specifically on reference or meaning for thinking that there are such exquisitely fine-tuned reference-determining factors. Second, that there are such factors seems extremely implausible given what we do understand about how reference is determined. For example, suppose Betty thinks "I hope my blind date is bald." According to Williamson, the occurrence of 'bald' in Betty's thought refers to some precise property such as the property of having fewer than 3,832 hairs on one's scalp, and this reference fact is fixed by substantial causal, social, or conceptual role properties of the token. Now, the reference of that mental occurrence of 'bald' can't be entirely fixed by the use of 'bald' in English, owing to the extreme context dependence of 'bald'. It's most plausible to suppose that its reference is determined by counterfactual use, or conceptual role, properties of it. This in turn seems to be fixed by the conditions under which 'My blind date is/isn't bald' would go into Betty's belief box. If that's so, then Williamson's view requires that however borderlinish the blind date might be, either 'My blind

date is bald' or 'My blind date isn't bald' will determinately go into Betty's belief box when he is visually present to her under conditions that are ideal for judging baldness. Yet that that isn't so is shown by the very fact that when Betty judges her blind date to be a borderline case of baldness, she believes neither that he's bald nor that he's not bald.

The third problem is closely related. We successfully communicate using vague language. In uttering 'I worked for a little while yesterday', I said, and you took me to say, that I worked for a little while yesterday. But it's extremely difficult to see how we could successfully communicate using vague language if Williamson's theory were correct. As regards the present example, it's very hard to see how the proposition I asserted could be the very same proposition you believed that I said. Suppose that my utterance of 'a little while' was such that the proposition I asserted is true just in case I did some work yesterday but worked no more than exactly two hours and forty-seven seconds (for some such truth condition must obtain if Williamson is right). The proposition I asserted is determined by the content of my causally operative contextual concept of a little time to work. The meaning of this concept, it's most reasonable to suppose, is determined, at least for the most part, by conceptual-role properties of the concept. Likewise, for your causally operative contextual concept of a little while to work. Yet it would seem extremely implausible that the conceptual role of your concept determined *exactly* the same span of time; perhaps your contextually relevant concept of a little time to work is any amount of time less than precisely three hours thirty-one and two-fifths seconds.

There are two ways Williamson might reply to this objection. First, he might argue that the facts are indeed as they seem: the proposition the speaker asserted is the same one the audience understood him to assert. To argue this, he will evidently have to argue that the contextually-variable reference of a vague term is determined by contextual factors equally available to speaker and audience. But on what basis can Williamson *argue* this as opposed merely to *claiming* it because his fragile position requires him to? Moreover, it's clear that such publicly available contextual factors don't generally determine the references of terms whose references are contextually variable. While it's true that a speaker relies on mutually known public factors to make known her referential intentions, it's her referential intentions, and not the contextual clues, which determine reference. If *S* says to *A* 'That book is profound', the reference of 'that book' is determined by *S*'s referential intentions, and the content of those intentions is even more clearly not determined by publicly accessible contextual factors.

In a published symposium on Williamson's *Vagueness*, I raised the successful communication problem, but without reliance on contextual variation of reference or meaning.²² In his reply,²³ Williamson cited the following passage from *Vagueness*:

Perhaps no two speakers of English match exactly in their dispositions to use 'thin'. It does not follow that no two speakers of English mean exactly the same by 'thin'.

For what individual speakers mean by a word can be parasitic on its meaning in a public language. The dispositions of all practitioners determine a sense that is available to each.²⁴

But this reply is of no help when the concern is with contextually-determined meaning and reference.

The second way Williamson might reply to the objection is to hold that it's virtually never the case that the proposition a speaker intends to convey is identical to the one his audience takes him to convey; there are always, or nearly always, slight differences, but differences that don't affect the purposes of communication. But this has some unappealing consequences. For one thing, it would follow that ascriptions of the form 'A said that (asserted, stated, etc.) that *S*' are virtually never true, since the proposition to which the saying-reporter is referring by the utterance of 'that *S*' isn't relevantly available to *A* for *A* to have said. Attempts to avoid this consequence are liable to lead to a version of super-valuationism wherein we deem a speaker to have said something true just in case each of an array of propositions, which includes the one she actually had in mind, are true.

Can an epistemic theorist who acknowledges that semantic properties are use dependent do better than Williamson with respect to not requiring miraculous, super-sensitive reference-fixing use-and-environment factors? Paul Horwich has given the impression that he thinks he can.²⁵ His account is a generalization of what he has to say about an artificial example he devises to illustrate the essential workings of his account:

Suppose we introduce a new term "glub" by means of the stipulation that integers greater than 20 are glub and integers less than 10 are not. What about 15? Here, as in the case of vague predicates, we can insist on classical logic, maintaining that 15 either is or is not glub and that there is a fact of the matter one way or the other, though it is impossible to know the truth... And in this case the explanation of our inevitable ignorance is fairly obvious. In order for the stipulated meaning of "glub" to cause an inclination to apply the word to 15, there would have to be an explanatory deduction from the fact that we apply "glub" in accordance with the stipulation, plus auxiliary facts, to the conclusion that we apply it to 15. But it is hard to see how there could be any such explanation...²⁶

Generalizing, Horwich tells us that *a* is a borderline case of *F* just in case one can't deduce from "the explanatorily fundamental regularity in our use of '*F*'" together with auxiliary facts that *a* is *F*, and that it's the failure to know in such a way which explains our inability to know that *a* is *F* when in fact *a* is *F*. "[T]he explanatorily fundamental regularity in our use of '*F*'," he tells us, "is approximated by a partial function $A(F)$ which specifies the subjective probability of its applying as a function of the underlying parameter *n* (i.e. "number of grains" for 'heap', "number of dollars" for 'rich', etc.)."²⁷

As I understand it, this attempt to explain the ignorance to which the epistemic theorist is committed is problematic.

(1) It's a consequence of Horwich's total view that the stipulation about 'glub' secures that 'glub' refers to the property of being glub. Now, for some number n between 9 and 21, the property of being glub = the property of being an integer greater than n . Let's say that the property of being glub = the property of being an integer greater than 17. Therefore, the stipulation about 'glub' secures that 'glub' refers to the property of being an integer greater than 17. So Horwich hasn't avoided the need for super-sensitive reference-determining factors; he's merely added another mystery to an already existing mystery: the mystery of how the stipulation about 'glub' can secure that the word refers to the property of being an integer greater than 17.

(2) Williamson has raised the following objection.²⁸ According to Horwich, we can know that 15 isn't glub if, and only if, we can deduce that proposition from the stipulation about 'glub' plus auxiliary facts. But among the auxiliary facts is, we're supposing, the fact that 15 isn't glub. Horwich can't amend his account by appeal to *known* auxiliary facts, for it's the unknowability of the fact that 15 isn't glub that he's seeking to explain.

(3) Horwich might amend his account by claiming that the stipulation has to occur *essentially* in the explanatory deduction; the proposition in question can't be deducible from the auxiliary propositions alone. Even as amended, however, the account remains problematic. It's a metaphysically necessary, language-independent fact that the property of being glub = the property of being an integer greater than 17. The stipulation about 'glub' simply secured that the new word referred to the property of being glub, i.e., the property of being an integer greater than 17. The metaphysically necessary, language-independent fact is unknowable, according to Horwich, but what explains the impossibility of knowing it? So what if we can't deduce that fact from the stipulation about 'glub' plus auxiliary facts, where the stipulation is essential to the deduction? If I stipulate that 'Lester Semester' shall name the number 17, my knowing that Lester Semester = $9 + 8$ isn't based, and evidently couldn't be based, on my having deduced that proposition from the stipulation plus auxiliary facts. All the premises in my knowledge-yielding deduction will be, and evidently must be, arithmetical.

There is, I believe, a line of reply available to Horwich, but it foregoes talk of explanatory deductions from fundamental regularities in use, and, in the end, threatens to undo the assumptions that lead one to be an epistemic theorist in the first place. First, he could take a deflationary line on the nature of properties and claim that although the property of being glub is language-independent in the sense that some numbers would have been glub however we used 'glub', still, the property is simply a hypostatization of the stipulation and its nature is determined by that stipulation.²⁹ Second, he could claim that the commitment to there being an identity of the form *the property of being glub = the property of being an integer greater than n* is simply an artifice of our commitment to bivalence. One can imagine Horwich being tempted to say "Look, there isn't *really* some number

that marks a sharp cutoff between being glub and not being glub; it's just that our commitment to bivalence requires us to say that there is." The problem with this line is that it makes it doubtful that we have the commitment to bivalence we're supposed to have. It makes us want to distinguish between being *really* true and being *so-to-speak* true, and that's a distinction that undoes both itself and the assumed commitment to bivalence which provokes it.

After having read these objections, Horwich, in a comment on a talk I gave,³⁰ claimed it's a mistake to suppose

that there must be some explanation, in terms of facts about how a given predicate is used, of why it stands for the particular property it does. For example, suppose being rich were constituted by having at least \$419,357. Then we are inclined to think that there would have to be some explanation, in terms of our practice with the word 'rich', of why it stands for *that* property (rather than, e.g., a property, richness*, constituted by having at least \$419,358).

But, Horwich argues, it's a mistake to think there would have to be such an explanation if 'rich' referred to the property of having at least \$419,357. He explains that:

An explanation, in terms of a predicate's use, of why it means what it does, can be expected only if there is a general reductive theory of meaning of the form

x means f -ness = $U(x, f)$

where $U(x, f)$ specifies the use of x in relation to objects of type f . But such a theory can be expected only if we take for granted that the reductive analysis of relational meaning-properties like

x means richness

must involve an analysis of the relation

x means y

And this conviction derives from the more general assumption that reductive analysis must preserve logical form (or, in other words, that if a property contains a given component, then whatever constitutes that property must either contain that same component or else must contain something that constitutes it). But this assumption about property-constitution is fallacious. A simple counterexample is that the property

x exemplifies flying

is constituted by the property

x flies

in which no analysis of 'x exemplifies y' can be found. Once we have talked ourselves out of the constitution fallacy, we will no longer demand a general theory of the form

$$x \text{ means } f\text{-ness} = U(x, f)$$

and therefore no longer expect to be able to explain why words mean what they do, given their use. And so it will be no objection to the supposition that a vague predicate has a sharp boundary that we are not able to explain why that boundary is precisely where it is.

It's important to appreciate that Horwich doesn't deny that meaning supervenes on use. There are, as we're about to see, epistemic theorists who do deny that, but Horwich isn't among them. For example, he wouldn't deny that 'rich' wouldn't refer to the property of being rich if we used the word the way we now use 'poor'. But if 'rich' refers to the property of being rich by virtue of its use (which we may take to include the social and physical environment in which it's used), then, by virtue of his acceptance of bivalence, Horwich must hold that there's a precise change in use which would secure that 'rich' no longer has its present meaning. This means there is a precise use fact on which the meaning of 'rich' currently supervenes. Let's call this fact *U*. What explains the fact that *U* determines 'rich' to refer to the property of having at least \$419,357? I take it Horwich wouldn't deny this dependency; rather, his claim would be that it can't be explained, and this because it could be explained only if we could have a reduction of the meaning relation which we can't have.

I'm not sure I see how this is supposed to help. Our initial puzzle with that version of the epistemic theory which recognizes semantic properties to be use dependent was our skepticism that there could be use properties that made it the case that 'rich' referred to the property of having at least \$419,357 as opposed to, say, the property of having at least \$419,357.01. Now we're told that while there are indeed such use properties, we can neither know what they are nor, if we could, explain why semantic properties of the kind in question supervene them. This seems to me to be the opposite of relieving mystery.

VI. Explaining the Ignorance: Use-Independent Semantic Properties

Extreme disquotationalists about semantic properties hold that meaning and reference are use-independent properties.³¹ According to them, instances of such schemas as

- 'S' is true iff *S*
- 'S' means that *S*
- 'F' is true of a thing iff it's *F*
- If *a* exists, then 'a' refers to *a*

are analytic, so that, for example, 'snow is white' would have been true iff snow is white even if English speakers had used 'snow' in the way they now use 'coal'.

When disquotationalism is conjoined with excluded middle and the standard understanding of falsity as truth of the negation, it follows that bivalence holds for sentences. So far, however, no consequences of interest can be drawn about vagueness, for to say that an utterance is true is just to say what the utterance says, and to say that an utterance is false is just to say what its negation says. The vagueness of ‘“Harry is bald” is true’ is on all fours with the vagueness of ‘Harry is bald’, and it’s consistent with this deflationary disquotationalist reading of ‘true’ that the proposition that Harry is bald (= the proposition that ‘Harry is bald’ is true) is itself neither true nor false. For this theorist to be an epistemic theorist it’s both necessary and sufficient for her also to claim that bivalence holds for propositions—the referents of that-clauses, the things we believe and assert. This should not be too big of a deal for the disquotationalist who subscribes to classical logic, since it merely requires accepting the instances of the schema

The proposition that *S* is true iff *S*,

or, more colloquially,

It’s true that *S* iff *S*.³²

There are, however, at least two problems with this way of being an epistemic theorist. First, there are all the problems endemic to holding that semantic properties such as reference, meaning, and truth are use independent. Not only is it reasonable to suppose that the fact that ‘Stephen Schiffer’ refers to me has something to do with how people use that name, but there are even hairier problems concerning the ascription of propositional attitudes to others and to the role of propositional attitudes in causal explanations. Second, while the disquotationalist epistemic theorist has the advantage of not having to account for miraculous reference-determining use-based factors, she still needs to explain why we can’t know either that Harry is bald or that he’s not bald, even though one of those propositions is true, or that baldness = the property of having fewer than 3,832 hairs on one’s scalp, and nothing in her disquotationalism, in and of itself, yields the needed explanation. If it’s a metaphysically necessary truth that whoever has fewer than 3,832 hairs on his or her scalp is bald, then why can’t we know this? Nothing about disquotational semantic properties answers this. Williamson has an account of this alleged ignorance, but the disquotationalist epistemic theorist can’t accept that account. What account can she offer?

I don’t see that Hartry Field has offered any published answer. If he has an answer, it must be in terms of his account of the definitely-operator. Field denies that the operator needs to be explained in truth-theoretic terms, or that it can be defined in any terms. He takes it to be a primitive operator that we come to understand in the same way we come to understand such operators as negation and disjunction and universal quantification: by learning to use it in accordance with certain rules. These rules would include the logical laws governing the op-

erator, and these would include a law that entailed that while ‘Definitely, Harry is bald or not bald’ must be true, ‘Harry is definitely bald or definitely not bald’ needn’t be true. Other rules of use for the operator tell us, for each vague term, what counts as the term’s definitely being true of a thing (it may do this in terms of resemblance to paradigms). Field continues:

There are also connections to our notions of knowledge and dependence on the physical facts: from ‘It is indefinite whether p ’ we can infer both ‘It is not knowable whether p ’ and ‘The physical facts don’t determine whether p ’; but the converse inferences are not acceptable (without a physicalist premise, in the second case, or a strong premise about our intellectual powers in the first). If it is indefinite whether p , there is not only no point in trying to find out, there is no point in even speculating: it is totally arbitrary what one says.³³

The point about there being no determination by physical facts seems to be either a mistake or else an allusion to some unexplained notion of determination, for Field wouldn’t deny that if Harry is a borderline case of a bald man, then so must be any other man whose hair situation was identical to Harry’s, and Field is committed to it’s being either true that Harry is bald or true that he’s not bald. A second problem with Field’s account is that his positive sketch of the definitely-operator precludes neither its being definable nor its being an epistemic notion (Williamson can accept Field’s gloss of the operator’s conceptual role—minus the mistake about physical determination, if it is a mistake). But the most important problem with Field’s conceptual-role account of the definitely-operator is that it patently does nothing to *explain* why we can’t know true borderline propositions.

Vann McGee and Brian McLaughlin are use-independent epistemic theorists who do attempt to answer the question. According to them, and *pace* Sorensen and Williamson, the definitely-operator isn’t an epistemic notion, and hence they’re free to give an account of it which does explain the ignorance to which epistemic theorists are committed. This is their strategy. According to them, we can’t know that borderline Harry is bald, when in fact he is, because while that proposition is true, it’s not *definitely* true. It’s incumbent on them, therefore, to give an account of the operator which supports this, and in “Distinctions without a Difference” they offer the following:

[D] x is *definitely* F iff (i) the thoughts and practices of the speakers of the language to which ‘ F ’ belongs determine conditions of application for ‘ F ’ and (ii) the facts about x determine that these conditions are met.³⁴

This is puzzling. Among the facts about Harry is the fact that he’s bald or the fact that he’s not bald. Suppose the former. How could the conditions of application of ‘bald’ fail to apply to a bald man? In “Timothy Williamson’s *Vagueness*” they respond to this worry. They in effect say that if they had a notion of a

precise fact, then they could revise [D] by inserting ‘precise’ so that (ii) now read ‘the *precise* facts about x determine that these conditions are not met’. But then they admit not to having a relevant notion of precise facts.

Now, if they had a notion of precise fact, condition (i) of [D] would be superfluous. Their real account would be:

x is *definitely F* iff it’s a precise fact that x is F .

Nor does the stuff in condition (i) about the thoughts and practices of speakers become relevant if we take them to be defining not what it is for a thing to be definitely such-and-such, but rather what it is for a sentence to be definitely true. This won’t help, for then their account should be:

‘ S ’ is definitely true iff it’s a precise fact that S .

All their stuff about thoughts and practices determining conditions for the application of vague terms is beside the point.

Finally, suppose we had a notion of precise fact. So what? Suppose it’s true that Harry is bald. How would the fact that that truth isn’t a “precise” fact explain why we can’t know it?

In a recent paper,³⁵ Brian McLaughlin has another go at his tenacious problem. He begins by distinguishing two relations holding between propositions: q *supervenes* on p and p *makes it the case* that q . The latter entails the former, but not vice versa. For example, the fact that $1 + 1 = 2$ supervenes on the fact that I have a nose, since there’s no possible world in which I have a nose and it’s not the case that $1 + 1 = 2$. But my having a nose doesn’t *make it the case* that $1 + 1 = 2$. McLaughlin makes no attempt to explicate the makes-it-the-case relation; he recognizes the desirability and ultimate need to say something about this, given the use he wants to make of the relation, but for now he simply takes it as primitive. The use to which he puts his distinction is as follows. First, he accepts bivalence for propositions, and thus the proposition that borderline Harry is bald is, for him, true or false. Second, he recognizes that every vague fact supervenes on some other fact. Suppose, for example, that borderline Harry is in fact bald. Then that fact supervenes on the exact hair situation on his scalp: you couldn’t have another person whose hair situation exactly matched Harry’s but who wasn’t bald. Let’s call this hair fact, on which Harry’s baldness supervenes, H . Third, and finally, it’s McLaughlin’s claim that while the fact that Harry is bald *supervenes* on H , it’s not the case that H *makes it the case* that Harry is bald. The proposal, then, is that

x is *indefinitely F* iff x is F , but nothing makes it the case that x is F .

The idea then is that what explains the fact that we can’t know borderline facts is that nothing makes them the case.

There are at least two problems with this solution. First, it's vacuous absent an account of the makes-it-the-case relation. If it's allowed that it's a fact that borderline Harry is bald, then *I* have no problem with saying that what makes it the case that he's bald is the hair situation on his scalp. Second, even if we accept that nothing makes indefinite facts the case, this doesn't yet explain why we can't know them. It's probably true that nothing makes it the case that there is something rather than nothing, but that hardly stops us from knowing it.

VII. A Concluding Suggestion

I hold out little hope for the epistemic theory of vagueness. My own views on these difficult issues are set out in a recent paper.³⁶ There I suggest that none of the familiar accounts of vagueness is right; they go wrong in attempting to give a happy-face solution to the sorites paradox, a solution that explains to us where and why sorites arguments go wrong in such a way that we won't ever again be perplexed by them. Vagueness, I propose, needs to be explained in terms of a special kind of partial belief, a kind of partial belief distinct from the familiar kind that is plausibly identified with subjective probability. Once vagueness is explained in this way, we'll see that it's indeterminate whether excluded middle, and thus bivalence, holds. Moreover, we'll see that the special kind of partial belief leads to paradox whether we adopt a bivalent or a non-bivalent notion of truth for propositions, and that it's because of this that the sorites must settle for an unhappy-face resolution, a resolution that shows the features of vagueness which lead to the paradox to be ineliminably endemic to the notion vagueness. I also propose that this failure to have a neat and tidy semantics and logic for vague notions comes at no great cost. Although it's indeterminate whether classical logic is true, it will continue to serve us well. In most cases we can harmlessly assume that the premises in an argument have truth-values and then apply classical logic with assurance of truth preservation. Our use of classical logic is strained when it confronts the sorites, but the sorites will strain us no matter what logic and semantics we use, and no important intellectual endeavor need suffer from this feature of our conceptual practices. Finally, this account of vagueness presupposes, and is supported by, a deflationary account of properties and propositions, whereby they are hypostatizations of certain trivial pleonastic transformations; in particular, the pleonastic transformation that takes us from

a is *F*

to its pleonastic equivalent

a has the property of being *F*

and the pleonastic transformation that takes us from any indicative sentence

S

to its pleonastic equivalent

It's true that *S*

which, we know, may be rewritten as

The proposition that *S* is true.³⁷

Needless to say, these are matters for another occasion.³⁸

Notes

1. Routledge, 1994.
2. See, e.g., Sorensen, *Blindspots* (OUP, 19??); Horwich, (Blackwell, 1990); Field, . "Deflationist Views of Meaning and Content," *Mind* 103 (1994): 249–85, and "Disquotational Truth and Factually Defective Discourse," *The Philosophical Review* 103 (July 1994): 405–52; McGee and McLaughlin, "Distinctions without a Difference," *The Southern Journal of Philosophy* 33 (1994): 203–51, and "Timothy Williamson's Vagueness," *Linguistics and Philosophy* 21 (April 1998): 221–235.
3. "Précis of *Vagueness*," *Philosophy and Phenomenological Research* 52 (1997): 921–928.
4. This may require qualification for the semantic paradoxes.
5. *Vagueness*, p. 187.
6. *Ibid.*
7. *Ibid.*
8. *Op. cit.*, p. 188.
9. See, e.g., Hartry Field, "Deflationist Views of Meaning and Content," *Mind* 103 (1994): 249–85, and "Disquotational Truth and Factually Defective Discourse," *The Philosophical Review* 103 (July 1994): 405–52.
10. Nor can one be ignorant of the truth-value of the *utterance* in a borderline case, which is to say that although one knows neither that *u* is true nor that *u* is false, there is no relevant proposition of which one is *ignorant*, even though *u* itself is either true or false. This is because neither *the proposition that u is true* nor *the proposition that u is false* is true, and this, of course, is because, on the view in question, if *u* expresses the proposition *P*, then the proposition that *u* is true = *P* and the proposition that *u* is false = the proposition that not-*P*.
11. "Vagueness, Truth and Logic," reprinted in R. Keefe and P. Smith, eds., *Vagueness: A Reader* (MIT Press 1996). The quote is from pp. 148–9.
12. "Distinctions without a Difference," *The Southern Journal of Philosophy* 33 (1994): 203–51; p. 217. I use superscripted asterisks as quasi-quotes.
13. This is consistent with the undeniable fact that the penumbras of vague terms dilate or constrict according to contextual interests. All that the point requires is that the that-clause in a particular utterance of, say, 'Sally said that Harry is bald' can determinately refer to a vague proposition.
14. For the notion of partial meaning, see Hartry Field, "Quine and the Correspondence Theory," *Philosophical Review* 83 (1974): 200–228.

15. Talk in this argument of reference to propositions needn't be a departure from the official ontologically neutral reading of that-clauses, for such talk may itself be understood in an ontologically neutral way.
16. Still with possible qualification with respect to the semantic paradoxes.
17. See *Blindspots*.
18. *Vagueness*, p. 215. See also his "Definiteness and Knowability" and "Reply to Commentators," *Philosophy and Phenomenological Research* 52 (December 1997): 945–953.
19. See his "Reply to Commentators."
20. "Contextualist Solutions to Scepticism," *Proceedings of the Aristotelian Society for 1995/6*: 317–333.
21. "Timothy Williamson's *Vagueness*."
22. "Williamson on Our Ignorance in Borderline Cases," *Philosophy and Phenomenological Research* 57 (December 1997): 937–943; the objection is raised on p. 942.
23. "Reply to Commentators," *Philosophy and Phenomenological Research* 57 (December 1997): 945–953.
24. P. 211.
25. "The Nature of Vagueness," *Philosophy and Phenomenological Research* 57 (December 1997): 929–935.
26. *Ibid.*, p. 932.
27. *Ibid.*, p. 933. I've changed Horwich's 'H' to 'F'.
28. "Reply to Commentators." Williamson also has a metalinguistic objection that strikes me as sound.
29. See my "Language-Independent Language-Created Entities," *Philosophical Topics* 24 (1996): 149–167.
30. At a conference in Oviedo, Spain, June 1998, sponsored by Sociedad Filosófica Ibero Americana.
31. See Hartry Field, *op. cit.*, and Vann McGee and Brian McLaughlin, *op. cit.* In fairness to these philosophers, they hold that our common semantic terms ambiguously express both use-independent and use-dependent terms. But this means they owe us *two distinct* solutions to the problems of vagueness, one for use-independent semantic notions and one for use-dependent semantic notions. Since I'm here concerned only with the former, we may harmlessly pretend that all relevant semantic notions are use independent.
32. Cf. Paul Horwich, *Truth* (Blackwell, 1990). As Horwich remarks, this may require some qualification in order to accommodate the semantic paradoxes.
33. *Op. cit.*, p. 420.
34. See p. 209.
35. "Supervenience, Vagueness, and Determination," *Philosophical Perspectives* 11 (1997): 209–230.
36. "Two Issues of Vagueness," *The Monist* 81 (April, 1998): 193–214.
37. This deflationary account of properties and propositions is elaborated in my "Language-Created Language-Independent Entities."
38. This paper has profited from comments on an earlier draft by Hartry Field, Kit Fine, Paul Horwich, Roy Sorensen, and Tim Williamson.