

Logic, Vagueness, and the Use Theory

STEVEN G. DANIEL
University of Nevada, Reno
Reno, NV 89557
USA

Introduction

In numerous essays over the years, and most recently in his *Meaning*, Paul Horwich has worked to articulate and defend a version of the use theory of meaning.¹ At the heart of his project is the idea that the meaning of a word is constituted by a regularity in speakers' *use* of it — i.e., by a regularity in speakers' dispositions to accept as true, to reject as false, or neither to accept nor to reject sentences containing it. A 'use regularity,' thus construed, is a dispositional acceptance property distinctive of a particular word and, derivatively, of sentences containing the word.² For example, the use regularity constituting the meaning of 'red' can be very roughly characterized as speakers' disposition to accept the application of 'red' to observed surfaces if and only if they are clearly red. Of course, the use regularity constituting the meaning of a word can be described in

1 Paul Horwich, 'Wittgenstein and Kripke on the Nature of Meaning,' *Mind and Language* 5 (1990) 105-21; 'Meaning, Use, and Truth,' *Mind* 104 (1995) 355-68; 'The Nature of Vagueness,' *Philosophy and Phenomenological Research* 58 (1997) 929-34; *Truth*, 2nd ed. (Oxford: Oxford University Press 1998); *Meaning* (Oxford: Oxford University Press 1998)

2 The correspondence between words and meaning-constituting use regularities will not always be one-to-one. Horwich allows that use regularities sometimes constitute the meanings of pairs or clusters of words, so that while each word in the cluster has a distinct meaning, those meanings are 'grounded' in the same use regularity. See Chapters 3 and 9 of his *Meaning* for details.

different ways; in particular, in describing a use regularity we might merely mention, and not use, the corresponding word. Furthermore, not every regularity in the use of a word is meaning-constituting. The use regularity that constitutes the meaning of the word 'tall,' for example, does so in virtue of the fact that, taken together with general psychological laws and the use regularities governing other words, it is *basic* to the explanation of speakers' overall pattern of use of 'tall.' Hence, a central claim of Horwich's use theory is that linguistic use is in general to be 'unified and explained in terms of a relatively small and simple body of factors and principles including, for each word, a basic use regularity' (*Meaning*, 45).

Horwich has also argued that vague languages are governed by the laws of classical logic. More specifically, he accepts what I will call the Classical View of vague languages: no departure from classical logic will make the phenomenon of vagueness more tractable and, in particular, we are to resist the familiar temptation to deny or otherwise 'weaken' our commitment to bivalence and/or excluded middle.³ Because Horwich accepts the Classical View, he does not take the vagueness characteristic of natural languages to affect the truth conditions of vague sentences or the extensions of vague predicates. Vagueness, he believes, is not a *semantic* phenomenon at all. In this respect, Horwich is in league with Timothy Williamson and the proponents of an epistemic theory of vagueness.⁴ However, Horwich is not an epistemicist. Although he allows that speakers suffer a form of necessary ignorance in borderline contexts, Horwich believes that the primary job of a theory of vagueness is to explain the peculiar phenomenology of vagueness; and the best explanation of the phenomenology of vagueness will, in his view, cite the 'paralysis of judgment' suffered by speakers in borderline cases, rather than speakers' epistemic limitations.⁵

3 In dubbing this the 'Classical View,' I do not mean to suggest that it is an historically popular view. It is not.

4 Timothy Williamson, *Vagueness* (New York: Routledge 1994); 'Definiteness and Knowability,' *Southern Journal of Philosophy* 33 (1995) 171-91; 'Wright on the Epistemic Conception of Vagueness,' *Analysis* 56 (1996) 39-45; 'Putnam on the Sorites Paradox,' *Philosophical Papers* 25 (1996) 47-56; 'What Makes it a Heap?' *Erkenntnis* 44 (1996) 327-9; 'Reply to Commentators,' *Philosophy and Phenomenological Research* 58 (1997) 945-53. See also R. Sorensen, 'An Argument for the Vagueness of "Vague"' *Analysis* 45 (1985) 134-7; 'The Vagueness of Knowledge,' *Canadian Journal of Philosophy* 17 (1987) 767-804; *Blindspots* (Oxford: Oxford University Press 1988); 'The Ambiguity of Vagueness and Precision,' *Pacific Philosophical Quarterly* 70 (1989) 174-83; 'Vagueness Within the Language of Thought,' *Philosophical Quarterly* 41 (1991) 389-413; 'A Thousand Clones,' *Mind* 103 (1994) 47-54.

I will not explore this dispute between Horwich and Williamson; for it is not their differences of opinion that interest me here, but rather their shared attraction to the Classical View and to a use theory of meaning. Furthermore, I will do no more than gesture at one source of the attraction: deflationism. Horwich embraces both a use theory of meaning and the Classical View of vague languages because he believes both to be motivated, if not entailed, by his deflationary (in Horwich's terms, 'minimalist') account of truth.⁶ While Williamson has not explicitly defended any particular brand of deflationism, he operates with what looks to be a deflationary conception of truth in his discussions of vagueness.⁷ And, like Horwich, Williamson has argued both that non-classical logics fail to yield satisfactory treatments of vagueness and that

-
- 5 Why, in Horwich's view, do borderline contexts induce a paralysis of judgment in speakers? Because the use regularities governing vague predicates possess 'a "gappy" character, specifying that the predicate is applied to objects possessing some underlying property to at least a certain specified degree, x , and that its negation is applied to objects possessing that property to less than a certain specified degree, y , where y is less than x , and that neither the predicate nor its negation is applied to objects possessing the property to some degree between x and y ' (*Meaning*, 64). In a borderline case, the relevant underlying property is present to a degree between x and y , which leads to a paralysis of judgment on the part of speakers. In Horwich's view, an explanation of vagueness in terms of gappy use regularities is preferable to an explanation in terms of an epistemic failing *external* to speakers (e.g., a reliabilistically construed ignorance of sharp semantic boundaries).
- 6 Horwich's own account of his commitment to the Classical View is in some respects curious. In his *Truth*, he claims that logic and the theory of truth have nothing to do with one another: '[Minimalism] is the proper conception of truth even in the context of deviant logics such as intuitionism or quantum logic, and would not be undermined by any arguments demonstrating the preferability of non-classical rules of inference' (74). This would seem to free Horwich of any commitment to the Classical View, since (e.g.) intuitionism denies the law of excluded middle. But Horwich then proceeds to argue that, given the equivalence schemata for truth and falsity (' $\langle p \rangle$ is true iff p ' and ' $\langle p \rangle$ is false iff not- p '), we cannot deny excluded middle except on pain of contradiction. In a clarificatory vein, Horwich notes that these results concerning excluded middle 'do not derive solely from the minimal theory of truth, but depend also on our having defined falsity as the absence of truth' (*Truth*, 77). But this remark clarifies little; for the definition of falsity as the absence of truth stems directly from the equivalence schemata which form the core of Horwich's minimalism, together with the minimalist claim that those schemata *exhaust* the concepts of truth and falsity.
- 7 As when he says, 'Given that an utterance says that TW is thin, what it takes for it to be true is just for TW to be thin, and what it takes for it to be false is for TW not to be thin. No more and no less is required. To put the condition for truth or falsity any higher or lower would be to misconceive the nature of truth or falsity' (*Vagueness*, 190).

the Classical View is consistent with a use theory of meaning. A clear explanation of these positions requires a discussion of deflationism. I cannot provide such a discussion here; for there are many different varieties of deflationism and a clear exposition of any one of them would take us too far afield.⁸ Aside from an occasional clarificatory remark, it will have to suffice to note that some writers hold a deflationary theory of truth to entail the Classical View and a more general deflationary program in semantics to require a use theory of meaning.⁹

Fortunately, my aim in this paper turns neither on what the proper formulation of deflationism is nor on what deflationism, properly formulated, entails. My aim is simply to show that a use theory of meaning cannot be reconciled with the Classical View of vague languages. More modestly, I will argue that, if one is a use theorist, then one had better be prepared to embrace a form of vagueness/indeterminacy which cannot be explained in terms of any type of cognitive limitation on the part of speakers.¹⁰ This, if correct, renders a use theory inconsistent with the positive accounts of vagueness offered by Horwich and Williamson. However, if we then proceed to make the common and, I think, plausible

8 For fear of leaving some readers high and dry, I offer the following very crude synopsis of Horwich's position. According to Horwich's version of deflationism, there can be no explanation of what the truth or falsity of a proposition consists in over and above what is stated in such equivalence schemata as ' $\langle p \rangle$ is true iff p ' and ' $\langle p \rangle$ is false iff not- p .' These schemata are supposed to be *exhaustive* of our concept of truth and they make no provision for departures from classical logic. Hence Horwich, *qua* deflationist, finds it plausible to deny that there are (e.g.) truth value gaps or failures of excluded middle.

As for the relationship between deflationism and the use theory of meaning, it is arguably impossible to be a deflationist about (e.g.) truth without also being a deflationist about other fundamental semantic concepts, such as meaning and reference. These concepts form a tightly woven fabric and are arguably interdefinable. For this reason, the deflationist about truth is apt to gravitate toward a deflationary approach to the theory of meaning, and in Horwich's view a use theory of meaning is the perfect vehicle for such an approach; for a use theory is consistent with the view that 'the basic use regularities of different words — like different laws of nature — need have no common form; and they need not relate the words they govern to the members of their extensions' (*Meaning*, 113).

9 For a good, general discussion of deflationary theories of truth and of deflationism in semantics more generally, see the introduction to S. Blackburn and K. Simmons, eds., *Truth* (Oxford: Oxford University Press 1999). Also, see chapter 4 of Horwich's *Meaning* for an explanation of how a deflationary theory of truth 'helps to dissolve a certain problem regarding aboutness — the notorious problem of intentionality — and thereby puts us in a position to discern the nature of meaning' (103).

10 Like Horwich and Williamson, I will use 'vague' and 'indeterminate' interchangeably.

assumption that an account of vagueness as a species of ignorance, a paralysis of judgment or as some other cognitive limitation is the only serious alternative to a formal semantic treatment of vagueness, then it follows that the use theorist had better be prepared to embrace some non-classical logic.¹¹

I will not argue that the Classical View of vague languages is false; nor will I argue that we should reject the use theory of meaning. My claim is simply that the two are inconsistent, and I leave it to the reader to decide how best to respond to the inconsistency. The dedicated use theorist is free to conclude that the problem lies with the Classical View (and with any varieties of deflationism which entail the Classical View). Deflationists, convinced that a departure from classical logic will create more problems than it will solve, are free to reject the use theory of meaning. My point is just that there is no avoiding these choices, and for some they promise to be difficult choices.¹²

I The Use Theory and the Classical View of Vague Languages

Horwich argues that a use theory of meaning is consistent with the Classical View of vague languages. In his view, the use theorist can consistently subscribe to the Classical View so long as he holds that, corresponding to every vague predicate, there is a use regularity that (i) is explanatorily basic with respect to speakers' overall deployment of the predicate, and that (ii) determines a *precise* meaning (hence extension) for the predicate.¹³

11 Of course, I am also assuming that any formal semantic treatment of vagueness will involve a non-classical logic. This too is a common and, in my view, plausible assumption; but I will not attempt to defend it here.

12 My discussion will focus almost exclusively on Horwich's version of the use theory. There are several reasons for this. First, there are many possible use theories in logical space, and I cannot address all of them. Second, Horwich's use theory captures, at least in my view, all of the historically essential features of a use theory of meaning. Third, Horwich's development of the use theory is probably the most sophisticated one currently to be had. Fourth, Horwich argues persuasively that his use theory *is* consistent with the Classical View. So, I focus on Horwich's use theory not because it is a frail specimen of the species but, on the contrary, because it is exceptionally robust.

13 Conditions (i) and (ii) are meant to hold so long as context is held fixed. The assumption of a fixed context will hold for all of the claims and examples I discuss in this paper, so I will not always make it explicit.

Horwich devotes considerable effort to explicating (ii); for it is not obvious how speakers' makeshift use of vague predicates could determine precise meanings for them. Patterns of use of vague predicates are often unprincipled, chaotic, and subject to myriad influences having nothing to do with linguistic use as such. One might doubt that precise meanings are explicable in terms of patterns of use such as these. However, it is no part of Horwich's *deflationary* use theory that speakers should be able to explain or 'read off' the precise meaning of a predicate, given knowledge of its corresponding use regularity. It is enough that the meanings of vague predicates be *determined* by, or that they *supervene* upon, underlying use regularities. Furthermore, if you believe that the use theory is obliged to tell an informative story about *how* a given meaning supervenes upon a given use regularity, then you are probably in the grip of some inflationary theory of truth or meaning — a condition Horwich will attempt to rectify. In fact, for Horwich it is precisely a deflationary approach to semantics that will relieve worries regarding the compatibility of the Classical View and a use theory of meaning; for a deflationary approach to semantics denies that a word's having precisely the meaning that it has is the sort of fact that need be, or even can be, explained. In Horwich's view, 'the basic use regularities of different words — like different laws of nature — need have no common form; and they need not relate the words they govern to the members of their extensions' (*Meaning*, 113).

This view of things is not implausible. A use theorist will of course hold that meaning supervenes upon use, and Horwich is right to point out that speakers might be unable to know precisely how meaning supervenes upon use. Nor is there any difficulty imagining that, in borderline contexts, the meanings of vague predicates (hence the truth conditions of vague sentences) supervene upon patterns of use in a way that is consistent with the Classical View, that respects bivalence, and so on. Supervenience relations such as these might seem unintuitive, but this in itself is no fatal objection. So, we should concede, I think, that the Classical View is consistent with the supervenience of meaning upon use. Yet this is to concede very little. The question before us is whether the Classical View is consistent with a use theory of meaning. Even if the Classical View is consistent with the supervenience of meaning upon use, it does not follow from this alone that the Classical View is consistent with a use theory. Before this latter consequence can be inferred, additional hurdles must be cleared.

The most imposing such hurdle derives from the fact that the use theory, simply in virtue of being a theory, purports to *explain* something. While the use theory will not explain why a particular word has precisely the meaning that it has, it does tell us that words in general mean what they do in virtue of the use regularities corresponding to them; and, more

importantly, it purports to explain the capacity of a use regularity to determine the meaning of a word in terms of that use regularity's being explanatorily basic with respect to speakers' overall deployment of the word. This means that, even after the use theory is harnessed to a deflationary agenda, there remain two different aspects of use regularities which, for the use theorist, must be perfectly coordinated: their role in the explanation of large scale patterns of linguistic use and their role in the determination of meanings. And if the use theorist holds that meanings are *precise*, as they are alleged to be on the Classical View, it is far from obvious how such coordination is to be achieved. Hence, there remains the worry that the Classical View might swamp the use theorist with more semantic precision than is intelligible from the point of view of a use theory of meaning. In the end, the Classical View will prove to be consistent with the use theory only if it is consistent with what a use theorist, *qua* use theorist, must say about the meanings of vague expressions, the truth conditions of vague sentences, the supervenience of meaning upon use, the explanatory role of use regularities ... in short, with the 'nuts and bolts' of a use theory.

We can certainly say something about how the use theorist might coordinate the meaning-constituting role of use regularities with their role in explaining broad patterns of linguistic use. Obviously, from the point of view of the use theorist, a word has meaning in virtue of the fact that speakers' use of it is characterized by a certain regularity. Also obvious is the fact that some regularities in use explain others. (For example, if, for some number n , English speakers are disposed to accept the application of 'heap' to n or more sand grains, this will explain their being disposed to assent to ' $n+10$ grains of sand make a heap.')

Moreover, if one use regularity explains another, it is reasonable to suppose that the former, rather than the latter, determines the meaning of the word; and should the former turn out itself to be explained in terms of some third use regularity, then this third use regularity would be the one that actually fixes the meaning of the word. And so on. In general, the use regularity playing the most important role in the explanation of speakers' overall deployment of a word is the use regularity that determines the meaning of the word. Motivating this picture is the commonsensical idea that '[o]ne of the properties of meaning that we recognize pretheoretically is that what people say is due, in part, to what they mean' (*Meaning*, 47).

Yet, in attempting to square a use theory of meaning with the Classical View, it is with the notion of an explanatorily 'basic' use regularity and condition (i) that we encounter the real source of problems. Why, for example, should we accept the claim at the very heart of Horwich's use theory: that corresponding to every word there is an *explanatorily basic* use regularity that constitutes its meaning? Talk of use regularities is

harmless enough, and it seems uncontentious to say that some use regularities are more important than others when it comes to explaining large-scale patterns of linguistic use. But the claim that the meaning of a word is always constituted by a use regularity that is explanatorily basic with respect to speakers' overall deployment of the word is a strong one indeed, and far from transparent. Here as elsewhere, it is not obvious what it means to speak of an explanation as 'the best' or as 'basic.'

Consider some mundane piece of behavior, such as my raising my right arm. There are, notoriously, a number of different explanations to be had of my raising my arm. Does one of them have to be 'the best' explanation? Perhaps there is a psychological explanation of my raising my arm, but must there be a single best psychological explanation? Must there be a single best physical explanation of the raising of my arm? One explanation might qualify as the most comprehensive, another might qualify as the most perspicuous; and perhaps two distinct explanations could rank as equally comprehensive and equally perspicuous. Is there any reason to suppose that there is some unique explanation to be singled out as 'basic'? None of these questions are easily answered, and there is no reason to believe that they will be more easily answered when the explanation of linguistic behavior is at issue.

In some respects it would be positively shocking if there were a unique, best explanation of speakers' overall deployment of a vague predicate. As mentioned earlier, the problem with speakers' overall patterns of linguistic use is not just that they are vast and sprawling, but that they are unprincipled, chaotic, and subject to myriad influences having nothing to do with linguistic use as such. Consider some vague predicate of English, *F*. Speakers will sometimes, through no fault of their own, falsely apply either *F* or its negation, and they will sometimes apply *F* or its negation when they should know better than to apply either. Speakers' use of *F* on particular occasions might be determined, not just by their (implicit) knowledge of associated use regularities, but also by feelings of impatience, a bad mood or a desire to resuscitate a stalled conversation. Speakers' use of *F* will change over time — sometimes resulting in a change in *F*'s meaning, sometimes not. Furthermore, speakers' use of *F* will often be inconsistent from one moment to the next. No matter how accurate and comprehensive our view of such patterns of use, any explanation of them in terms of discrete use regularities will involve a good measure of *idealization*, and there need be no reason to prefer one idealization to all others; for one idealization might yield explanations of use that are neither better nor worse than those yielded by another.

II The 'Heap Case'

An example, which I call the Heap Case, brings my worries into sharp focus.

The Heap Case: Suppose that, for some number n , 50% of the speakers in a given speech community are disposed to accept the application of 'heap' to n grains of sand while the remaining 50% of speakers are disposed to reject the application of 'heap' to n grains of sand. Let us further suppose that, for any number m other than n , speakers are unanimous in either accepting, rejecting, or in neither accepting nor rejecting the application of 'heap' to m grains of sand. On the whole, speakers' use of 'heap' in this community is similar to our own in its outward appearances. There are collections of sand grains which speakers recognize to be clear instances of heaps, and other collections of sand grains recognized to be borderline heaps. We can even imagine that this community's classification of collections of sand grains as clear heaps and borderline heaps is, on the whole, recognizably similar to our own. Furthermore, just as our judgments about what is and is not a heap grow noticeably less confident in borderline contexts, so too, for some numbers ($k_1 \dots k_n$) (including n), speakers in this community are noticeably less confident in accepting, rejecting or in neither accepting nor rejecting the application of 'heap' to ($k_1 \dots k_n$) grains of sand.

The Heap Case is a description of one speech community's pattern of use of 'heap' at a particular instant in time. The pattern of use of 'heap' it describes is obviously possible. No doubt it is highly unlikely (and highly unstable), but then so is any other pattern of use described at this level of detail.

Opinions may vary as to whether, in the Heap Case, n grains of sand constitute a borderline case for '... is a heap.' Some writers characterize borderline cases in terms of conflicting patterns of use, and the Heap Case illustrates a conflict in speakers' pattern of use of 'heap.' Yet other writers prefer to characterize borderline cases in terms of speakers' failure either to accept or to reject the application of a vague predicate.¹⁴

14 Of course, to prefer one characterization of 'borderline case' is not necessarily to oppose others. For example, while Williamson typically describes borderline cases in terms of indecisive rather than conflicting patterns of use, he does not seem hostile to other characterizations. This is plain when he says, 'Some writers on vagueness, such as C.S. Peirce and Crispin Wright, characterize borderline cases in terms of the conflict rather than absence of opinion. Neither is essential; what matters is the absence of knowledge' ('Reply,' 946).

If this is how we understand 'borderline case,' then n grains of sand are not borderline for '...is a heap' in the Heap Case. That some writers prefer one of these alternative characterizations of 'borderline case' to the other is interesting and important; yet for my purposes it does not matter whether we regard n grains of sand as a borderline heap. All that matters is that 'heap' is vague in the community I have profiled and that the pattern of use of 'heap' I have described is possible. Both of these facts I take to be beyond serious doubt.

In that case, how might the use theorist best explain, in terms of a unique use regularity corresponding to 'heap,' the pattern of use described in the Heap Case? There are three obvious possibilities — i.e. three candidate use regularities — to consider. One candidate use regularity, a , licenses the application of 'heap' to n grains of sand; another use regularity, b , licenses the application of 'non-heap'; a third use regularity, c , licenses the application of neither 'heap' nor 'non-heap.' Since, for any number m other than n , speakers are unanimous in either accepting, rejecting, or in neither accepting nor rejecting the application of 'heap' to m grains of sand, use regularities a , b and c are in perfect agreement for all numbers ($m_1 \dots m_n$). a , b and c differ only about whether 'heap' is to be applied to n grains of sand.

Which use regularity provides the best explanation of speakers' use of 'heap'? There is clearly nothing to recommend a over b , or vice versa, since we can view each as explaining the linguistic use of half the population (while failing to explain that of the other half). This leaves us with c ; but c is the worst performer of all, since no speakers are disposed to remain silent about whether n grains of sand make a heap. Does either a or b have to constitute a better explanation in the Heap Case? Obviously not, since they perform equally well (or poorly). And no matter how we poke and prod the speakers in this community to uncover their dispositions to use 'heap,' there is *obviously* no reason why our investigation (or, indeed, the facts themselves) *must*, in the end, favor either a or b .

One might object that my presentation of the Heap Case is simplistic on the grounds that a , b and c are not necessarily the only use regularities available to the use theorist. Following Horwich, I have described use regularities in relatively simple terms as dispositional 'acceptance properties' of words. However, such properties need not be as simple as a , b and c , and a more detailed description of the Heap Case might reveal additional use regularities of great value to the use theorist. Suppose, for example, that fewer than half the speakers in my imagined community are disposed to accept some logical compound of which ' n grains of sand make a heap' is a constituent. The use theorist might argue that this and other similar facts about use could spoil the symmetry at the heart of the Heap Case.

It goes without saying that a different, more detailed description of the Heap Case might require an analysis very different from the one provided above. However, this fact implies nothing about how the use theorist should address the Heap Case as I have presented it; nor is there any dialectical reason to regard my Heap Case as in need of further elaboration. While a more detailed description of the scenario might reveal new, explanatorily potent use regularities to the use theorist, it certainly *need not* reveal such regularities. The addition of further detail might simply leave the use theorist facing some more complex explanatory impasse.¹⁵ Indeed, the real difficulty facing the use theorist's attempts to explain symmetrical and conflicted patterns of use of the kind profiled in the Heap Case can only be understood in relation to the idealizing nature of explanation. This feature of explanation is completely general. It is tied neither to a particular way of describing use regularities nor to a particular approach to the weighing of explanations, and it manifests itself in many different areas of inquiry. So, how the use theorist describes his use regularities is not a dire concern. Variations of the Heap Case are easily engineered; the possibility of competing yet equally matched use theoretic explanations remains a fact of life.

In a related vein, one might object that the Heap Case appears problematic for the use theorist only because it harbors an objectionable, inflationist conception of meaning. If any one of the use regularities featured in the Heap Case were explanatorily basic, we would be able to read the extension of '... is a heap' off of it; yet use regularities need not enable us to do this. Hence, there could be an explanatorily basic use regularity present in the Heap Case, only operating at a different, less 'semantically transparent' level of description.

This objection again underestimates the force of the Heap Case. The charge that *a*, *b*, and *c* are inflationary is particularly weak in connection with *c*: since *c* licenses the application of neither 'heap' nor 'non-heap' to *n* grains of sand, how could we read the truth conditions of '*n* grains make a heap' off of it? Yet my description of use regularities *a* and *b* also implies nothing about how the truth conditions of '*n* grains make a heap'

15 Suppose someone suggested that, if only we employed a richer, more sophisticated psychological idiom, then we would see that there is a unique, best psychological explanation of every instance of human behavior. While most everyone allows that such an idiom would sometimes assist us in the task of explaining behavior, few would find the suggestion palatable. For taking up a more sophisticated idiom might be no help at all and might make matters worse by greatly multiplying the number of equally matched explanations from which we are to choose. Such are the vagaries of explanation, including not only psychological but also *use theoretic* explanation.

(or of any other vague sentence) might supervene upon them. Even if use regularity *a* (or *b*) were explanatorily basic, this would imply only that the truth conditions of '*n* grains of sand make a heap,' whatever those happen to be, supervene upon *a* (or *b*). It would *not* imply that '*n* grains of sand make a heap' is true (or, in the case of *b*, false). So, there is nothing inflationary about the Heap Case. It is fair to say that the Heap Case is somewhat artificial; for it operates with a stipulated pattern of use of 'heap' against which we can assess the explanations yielded by *a*, *b* and *c*. This is a luxury we do not enjoy in the field, where speakers' use of vague expressions is, for the most part, beyond our view. Yet this does not weaken the force of the Heap Case. The pattern of use of 'heap' described there is perfectly possible and so the use theorist should be prepared to address it.¹⁶

III Implications of the Heap Case

The Heap Case poses a problem for anyone who wishes to reconcile a use theory of meaning with the Classical View of vague languages. To see why, we must attend to three things. Notice, first of all, that while use regularities *a* and *b* yield competing explanations of speakers' use of 'heap,' we nevertheless feel confident asserting that, in the speech community profiled in the Heap Case, neither *a* nor *b* is explanatorily basic in the sense of providing the unique, best explanation of speakers' use of 'heap.' Hence the use theorist can accept neither '*a* best explains speakers' use of "heap"' nor '*b* best explains speakers' use of "heap"' as true; for given that neither *a* nor *b* provides a better (or worse) explanation of speakers' use of 'heap' than the other, the truth of either of the aforementioned sentences would be simply *unintelligible* to the use theorist. The only facts in virtue of which one of these sentences could be true are *a*'s being more explanatorily potent than *b*, or vice versa, and neither of these facts obtain.¹⁷

16 My discussion of the Heap Case was improved by comments from two anonymous referees.

17 There is no reason to complain that my invocation of 'facts' here is inflationary. I do not have a theoretically 'loaded' notion of fact in mind, and the deflationist is perfectly at home with the idea that true sentences (or propositions, if you prefer) 'owe their truth' to the facts. It is even consistent with deflationism to say that 'Grass is green' is true *because* grass is green, or to say that 'Grass is green' is false *because* grass is not green. What the deflationist denies is the philosophical thesis that the truth of a sentence is in general *constituted* by its bearing some well-defined, explanatory relation to 'the facts.' Horwich makes this point by remarking that his

The second thing to notice is that the use theorist cannot reject both ‘*a* best explains speakers’ use of ‘heap’ and ‘*b* best explains speakers’ use of ‘heap’ as false. The possibility that both of these sentences are false is at least intelligible to the use theorist. However, in the use theorist’s view, the meanings of vague expressions supervene upon explanatorily basic use regularities, and the explanations yielded by *a* and *b* are the best explanations to be had of speakers’ use of ‘heap.’ Hence, deeming both of the above sentences false would entail that ‘heap’ is meaningless; yet ‘heap’ is clearly as meaningful in this community as in our own. The use theorist is not a meaning skeptic. In that case, he cannot reject both of the above sentences as false, and this leaves him with no alternative but to hold that the meaning of ‘heap’ supervenes upon a use regularity that is, at least in some sense, *indeterminate* as between use regularity *a* and use regularity *b*.

Given that the Heap Case presents us with a form of indeterminacy, we must inquire as to the nature of that indeterminacy. In particular, is the indeterminacy explicable in terms of speakers’ ignorance or in terms of some other cognitive limitation? This brings us to the third thing to notice about the Heap Case. The indeterminacy presented by the Heap Case affects the *facts* about explanation. It is not the case that either *a* or *b* provides the best explanation of speakers’ use of ‘heap,’ only we are constitutionally unable to know which of them does so. A particular use regularity certainly *might* provide a better (or worse) explanation of large-scale patterns of linguistic use than we are in a position to recognize. So much is obvious, given what Williamson (*Vagueness*) calls the ‘sprawling quality’ of linguistic use. However, *all* of the facts relevant to explanation are made *explicit* in the Heap Case; they are not hidden. In the Heap Case, every individual speaker’s disposition to accept, reject or neither to accept nor reject the application of ‘heap’ to *n* grains of sand is taken into account. Yet, for the use theorist, the indeterminacy affecting *a* and *b* *persists*.

It is useful at this point to introduce a distinction between strong indeterminacy and weak indeterminacy. I will use ‘strong indeterminacy’ to refer to forms of vagueness/indeterminacy which cannot be explained in terms of the cognitive limitations of speakers, and I shall continue to assume that any form of vagueness/indeterminacy which cannot be thus explained requires a formal semantic treatment and some

minimalism ‘does not deny that truths *do* correspond — in *some* sense — to the facts; it acknowledges that statements owe their truth to the nature of reality; and it does not dispute the existence of relationships between truth, reference, and predicate satisfaction’ (*Truth*, 105).

revision of classical logic.¹⁸ I will use 'weak indeterminacy' to refer to forms of vagueness/indeterminacy which are amenable to classical logic in virtue of their being explicable in terms of the cognitive limitations of speakers. If it is weakly indeterminate that p , then p is either true or false, but we are unable to know that p is true (or false). Corresponding to both forms of indeterminacy is a corresponding form of determinacy; and, obviously, its being strongly determinate that p is consistent with its being weakly indeterminate that p .

In terms of this distinction, the Heap Case implies that, from the point of view of the use theorist, there is form of strong indeterminacy affecting the truth conditions (hence the meanings) of the sentences ' a best explains speakers' use of "heap"' and ' b best explains speakers use of "heap."' Furthermore, in the Heap Case the use theorist can say, of use regularity a , neither that it best explains speakers' use of 'heap' nor that it does not best explain speakers' use of 'heap.' (And so too for use regularity b .) The use theorist is thus encouraged to deny or, at any rate, somehow 'weaken' his commitment to either bivalence or excluded middle.¹⁹ Of course, so far as the proponents of the Classical View are concerned, it is pointless to break with classical logic. This may well be true. Yet what the Heap Case is designed to illustrate is simply that the Classical View is *unintelligible* from the point of view of a use theory of meaning. In the Heap Case, the use theorist must either embrace strong indeterminacy or remain silent; hence the proponents of the Classical View must reject the use theory.²⁰

18 The details of strong indeterminacy I am happy to leave open. If it is strongly indeterminate that p , this might be taken to mean that p is neither true nor false; yet strong indeterminacy might be taken to be a primitive notion on a par with truth and falsity. Thus, its being strongly indeterminate that p might be taken to mean that there is simply no true answer to the question whether p is true (or false) that does not invoke the notion of strong indeterminacy, *including* the answer that p is neither true nor false. Strong indeterminacy can be construed in other ways as well, but this is not my concern here. Indeed, it is consistent with the argument of this paper that the notion of strong indeterminacy is ultimately *unintelligible*. My claim is merely that it is a notion indispensable to the use theorist.

19 In other words, the use theorist is forced to fall back either on a system of logic that denies bivalence and/or excluded middle or on one that fails to assert one or both of these principles. For example, the use theorist might at this point be attracted to supervaluationism, which denies bivalence while retaining excluded middle. Of course, he must then confront Williamson's (*Vagueness*) arguments to the effect that supervaluationism is unable to accommodate the phenomenon of higher-order vagueness.

20 Notice that if, like Horwich, we accept the claim that meaning is *compositional*, then

IV Replies to the Heap Case

It is up to the proponent of the Classical View (the 'classical theorist') to find some flaw in my analysis of the Heap Case. In particular, he must show that what appears to be the strong indeterminacy affecting use regularities in the Heap Case is in fact merely weak indeterminacy. The argument most frequently used to this end by the classical theorist involves an appeal to ignorance or to some other type of cognitive limitation. I believe I have already shown that no such argument can succeed. There is nothing 'hidden' from us in the Heap Case. Hence, in order to undermine my analysis of the Heap Case, the classical theorist's strategy must focus on my characterization of the use theory, the use theorist's ability to coordinate the meaning-constituting role of use regularities with their role in explaining large scale patterns of linguistic use, and the various ways in which meaning might be thought to supervene upon use.

All the classical theorist must show is that a use theory of meaning is *consistent* with the Classical View, and there are a few considerations that might help in this regard. First, there is some sense in which '*a* best explains speakers' use of "heap"' and '*b* best explains speakers' use of "heap"' are *both* true. While neither *a* nor *b* provides the unique, best explanation of speakers' use of 'heap,' both *a* and *b* yield explanations that are *no worse* than those yielded by any other use regularities. *a* and *b* thus have at least some sort of explanatory pedigree. Second, there is

my analysis of the Heap Case quickly turns into an argument for a completely general thesis of strong meaning indeterminacy. If meaning is compositional, then, insofar as the Heap Case shows that the meaning of '...best explains speakers' use of "heap"' is strongly indeterminate, it shows that the meanings of that predicate's constituent expressions are strongly indeterminate. And insofar as the meanings of 'best,' 'explain,' 'speaker,' and the like are strongly indeterminate, then so are the meanings of all of the sentences they can enter into, and the meanings of all the words figuring in those sentences, and so on until strong meaning indeterminacy is pervasive. Due to compositionality, strong indeterminacy is like a spot of dye in a pool of water: what begins as a small drop quickly spreads. The only direct way for the use theorist to resist this argument is by rejecting the compositionality of meaning. I doubt the use theorist will rejoice over this option; Horwich's use theorist certainly will not. Yet even if compositionality is rejected, there is still the strong indeterminacy affecting the meaning of '...best explains speakers' use of "heap"' in the Heap Case; and if we can engineer the Heap Case, we just as easily engineer the Red Case, the Bald Case, the Thin Case, the Rich Case, the Tall Case, and so on for any number of vague predicates. Even these isolated instances of strong meaning indeterminacy are inconsistent with the Classical View, and the use theorist cannot avoid them.

the fact that the classical theorist can tell any story he likes about the supervenience of meaning upon use in order to make the Classical View come out true — so long as that story is consistent with the use theory.

Since the problem posed by the Heap Case stems from the fact that *a* and *b* provide equally good, yet *competing*, explanations of speakers' use of 'heap,' the obvious feature of the Heap Case to question, I think, is its assumption that *a* and *b* are in 'competition' with one another for the explanation of speakers' use of 'heap.' There are at least two separate strategies available to the classical theorist for undermining this assumption. The first I will call the Fusion Argument.

At the heart of the Fusion Argument is the idea that the meaning of 'heap' has, as its 'supervenience base,' a use regularity. Since neither *a* nor *b* provides the unique, best explanation of speakers' use of 'heap,' the use theorist cannot say that the meaning of 'heap' supervenes upon either *a* or *b*. However, *a* and *b*, taken together, provide the best explanations there are, in terms of determinate use regularities, of speakers' use of 'heap.' So, let the use theorist simply say that the meaning of 'heap' supervenes upon both *a* and *b*, and let him view *a* and *b* as jointly explaining speakers' use of 'heap' — i.e., as jointly constituting a single, determinate use regularity. The resulting use regularity (*a, b*) provides the best explanation of speakers' use of 'heap' there is to be had, and it also constitutes a precise meaning for 'heap.' Thus, the Fusion Argument concludes, strong indeterminacy can be avoided in the Heap Case, both as it affects the use regularities *a* and *b* and the meaning of '...best explains speakers' use of "heap."'

The goal of the Fusion Argument is to coordinate the meaning-constituting role of use regularities *a* and *b* with their role in explaining large scale patterns of linguistic use, and to do so in such a way that respects the Classical View. Yet it is fatally flawed; for no sense can be made of the idea that *a* and *b* can somehow 'fuse' to form another use regularity (*a, b*). While use regularity *a* bids speakers to apply 'heap' to *n* grains of sand, use regularity *b* bids them to reject the application of 'heap' to *n* grains of sand. If we can imagine there to be a use regularity (*a, b*), then it is a use regularity which bids speakers simultaneously to accept and to reject the application of 'heap' to *n* grains of sand. In all likelihood the use theorist will not countenance such a use regularity. Yet even if he does, he gains nothing; for even if the use theorist views (*a, b*) as constituting a single, strongly determinate use regularity, this use regularity will yield *worse* explanations of speakers' overall use of 'heap' in the Heap Case than either *a* or *b* taken individually. None of the speakers in the Heap case both accept *and* reject the application of 'heap' to *n* grains of sand; hence the use regularity (*a, b*) fails to be explanatorily basic. In that case, it is of no interest to the use theorist.

Since *a* and *b* cannot combine to form a single, determinate use regularity, the use theorist must pursue a different line of argument. One such argument I will call the Overdetermination Argument. At the heart of the Overdetermination Argument is the intuition that, if *a* and *b* tie at providing the unique, best explanation of speakers' use of 'heap,' this can only be *because* they determine the same meaning for 'heap.' Of course, neither *a* nor *b* can determine a meaning for 'heap' in the speech community at large, since neither is explanatorily basic. However, it might be argued that meanings can supervene not only upon use regularities, but also upon a *collection* of use regularities, so long as each use regularity in the collection is such that, were it explanatorily basic, it would determine the same meaning for 'heap' as any other member of the collection, were any other member explanatorily basic. Thus, while the meaning of 'heap' is univocal in the speech community profiled in the Heap Case, that community consists of two distinct subpopulations, A and B. *a* provides the unique, best explanation of A's use of 'heap,' and *b* provides the unique, best explanation of B's use of 'heap.' Furthermore, *a* determines the same meaning for 'heap' in A that *b* determines for 'heap' in B.

On this picture, *a* and *b* are not in competition to provide the unique, best explanation of the use of 'heap' in the speech community. Since *a* and *b* determine the same meaning for 'heap' in A and B, there is no need for competition between them. In that case, what the meaning of 'heap' supervenes upon in the speech community is not a use regularity at all but a pair of distinct use regularities, *a* and *b*; in this sense, the meaning of 'heap' is 'overdetermined' in the speech community. While '*a* best explains speakers' use of "heap"' and '*b* best explains speakers' use of "heap"' are both false at the level of the speech community, '(*a* and *b*) best explain speakers' use of "heap"' is true insofar as *a* and *b* yield the unique, best explanations of A's and B's use of 'heap.'²¹ The conclusion then is that strong indeterminacy can be avoided in the Heap Case.

21 One might object that the Overdetermination Argument relies upon a defective notion of a speech community. If two populations just happen to use 'heap' with the same meaning, this does not imply that they form a single speech community. Indeed, if two isolated populations turn out, by happy coincidence, to speak English, this does not imply that they form a single community of English speakers. Speech communities are not merely communities in which a fixed set of words possess the same meaning, but communities in which the meanings of those words are fixed by the same use regularities. Perhaps this is so. In that case, what the Overdetermination Argument boils down to is the claim that my Heap Case presents us not with a single speech community that is divided, at the level of use, into two distinct subpopulations, but with different speech communities that differ subtly at the level of use.

The problem with this argument is visible in the intuition underlying it: that if *a* and *b* tie at providing the unique, best explanation of speakers' use of 'heap,' this can only be because they each determine the same meaning for 'heap.' Suppose that *a* and *b* do determine the same meaning for 'heap' in subpopulations A and B. In that case, we must ask why *a* (and *not b*) is explanatorily basic in A and *b* (and *not a*) is explanatorily basic in B; and there is no plausible use theoretic answer to these questions. This is unacceptable. 'Sameness of meaning' is something the use theorist purports to explain, whether it manifests itself in the guise of synonymy, univocality, or otherwise. So far as the use theorist is concerned, if *a* and *b* determine the same meaning for 'heap' whenever they are explanatorily basic, then *a* and *b* ought *always* to yield equally good (or bad) explanations of speakers' use of 'heap.'²² To deny this would be to set the meaning-constituting role of use regularities floating unacceptably free of their role in explaining large-scale patterns of linguistic use. However, we can obviously imagine a speech community in which *a*, and not *b*, (or *b*, and not *a*) is explanatorily basic. Thus, the intuition that *a* and *b* can tie at providing the unique, best explanation of speakers' use of 'heap' only if they determine the same meaning for 'heap' is undermined by the fact that, for the use theorist, *a* (and not *b*) could be explanatorily basic only if *a* is such that it would determine a *different* meaning for 'heap' than *b*, were *b* explanatorily basic.

This leaves the use theorist right where he was following my analysis of the Heap Case. He remains committed to the view that, since *a* and *b* provide equally good explanations of speakers' use of 'heap,' the use regularity constituting the meaning of 'heap' is *strongly indeterminate* as between *a* and *b*, and the meaning of 'heap' is in turn strongly indeterminate as between the meaning it would have if *a* were explanatorily basic and the meaning it would have if *b* were explanatorily basic.

V Horwich on Meaning Indeterminacy

To my knowledge, Horwich has not discussed the kind of problem created by the Heap Case. However, he does sometimes exhibit general concern about how the meaning-constituting role of use regularities is

22 Alternatively, suppose that *a* and *b* are such that, if *a* best explained speakers' use of 'heap' in some counterfactual situation, then *a* would determine a *different* meaning for 'heap' than *b*, were *b* to provide the best explanation of speakers' use of 'heap' in some other counterfactual situation. In that case, if the meaning of 'heap' is always *precise*, then *a* and *b* should *never* provide equally good (or bad) explanations of speakers' use of 'heap.'

to be coordinated with their role in the explanation of large-scale patterns of use. His concern manifests itself most explicitly in his discussions of the analytic/synthetic distinction. This is not surprising, since in asking whether there is, corresponding to each word, a unique use regularity that is at once meaning-constituting and explanatorily basic, one thing we are asking is whether there is a sharp distinction to be drawn between those regularities in our use of a word that are meaning-constituting and those that are not. If the use theorist insists that the meaning of every word is constituted by some unique, explanatorily basic use regularity, then he commits himself both to an analytic/synthetic distinction and to the view that meanings are subject to no very serious form of indeterminacy. On the other hand, if the use theorist does not wish to claim that the meaning of every word is constituted by some unique, explanatorily basic use regularity, then he is essentially rejecting an analytic/synthetic distinction and also opening the door to a very serious form of strong meaning indeterminacy. Such a use theorist could still hold that some use regularities are *more important than others* when it comes to explaining large scale patterns of linguistic use and to the determination of meanings; but he need not stake everything on the claim that, corresponding to every vague predicate, there is a unique, explanatorily basic use regularity.

There is no doubt that Horwich accepts the existence of some form of 'meaning indeterminacy.' Horwich says explicitly that his use theory 'implies that there may be words such that it is indeterminate whether or not a given meaning property is exemplified by them' (*Meaning*, 64). Furthermore, Horwich notes that

There will sometimes be alternative, equally good ways of finding a simple regularity in the use of a word that (in conjunction with the use regularities of other terms and with general psychological laws) will account for all aspects of its use. Therefore there will sometimes be no objective fact of the matter as to where the boundary lies between the pattern of use that constitutes the meaning of an expression, and other facts about its deployment. (*Meaning*, 60)

All Horwich *appears* to do here is to embrace a kind of indeterminacy affecting use regularities and to reject an analytic/synthetic distinction; and Horwich immediately proceeds to describe himself as conceding 'indeterminacy in the constitution of meaning' (*Meaning*, 60). But is the indeterminacy Horwich here embraces the kind of indeterminacy encountered in my Heap Case? Or, alternatively, is Horwich really rejecting an analytic/synthetic distinction?

If one surveys the spectrum of opinion about the analytic/synthetic distinction, Horwich can appear to occupy every position on it. In the passage just quoted, he seems to reject it. Elsewhere he describes his attitude toward analytic truth as 'equivocal, although more sympathetic

than not' (*Meaning*, 152), and his rejection of an analytic/synthetic distinction is often hedged, as when he says, 'Thus there is indeed no *useful* distinction between the analytic and the synthetic' (*Meaning*, 152; my emphasis). Yet Horwich clearly seems committed to an analytic/synthetic distinction of *some* kind, given that his use theory of meaning is founded on the idea that we may 'divide the entire usage of a word into those patterns (or regularities) of its use that are *explanatorily basic* ... and that we identify the explanatorily basic patterns with those that provide the meaning of the word' (*Meaning*, 151).

One reason why Horwich's attitude toward an analytic/synthetic distinction is difficult to pin down is that there are several different conceptions of analyticity, and Horwich is more sympathetic with some than others. Another reason is the familiar ambiguity infecting such words and phrases as 'indeterminacy' and 'no objective fact of the matter.' In terms of the distinction introduced earlier, indeterminacy comes in two varieties: weak and strong. So far as Horwich is concerned, the 'indeterminacy in the constitution of meaning' described above can only be weak indeterminacy. Given his allegiance to the Classical View and to the deflationary origins of the Classical View, Horwich must be read as embracing a strongly determinate analytic/synthetic distinction while rejecting only the claim that the distinction is weakly determinate. To read Horwich otherwise is to court disaster; for without a strongly determinate boundary between explanatorily basic use regularities and other, non-basic ones, Horwich would be forced to embrace strong indeterminacy at the level of use and, with it, strong meaning indeterminacy. Yet Horwich certainly does not embrace the existence of strong indeterminacy of any kind. In his view, use regularities are subject only to weak indeterminacy.

Notice, though, that the indeterminacy generated by the Heap Case is clearly not weak indeterminacy. In the Heap Case there is no reason to believe, and every reason to deny, that either use regularity *a* or use regularity *b* must provide the unique, best explanation of speakers' use of 'heap,' only we are constitutionally unable to judge which one does so.

More importantly, we are now in a position to say something about what the strong indeterminacy affecting use regularities *is*. In fact, I mentioned the source of the indeterminacy when I earlier observed that, no matter how accurate and comprehensive our view of a speech community's pattern of use of a word, any explanation of it in terms of discrete use regularities will involve a good measure of *idealization*, and there need be no reason to prefer one idealization to all others. The strong indeterminacy presented by the Heap Case, far from being a species of ignorance or the reflection of some cognitive limitation, is simply a by-product of the kind of idealization involved in explaining broad,

often messy, patterns of linguistic use in terms of discrete, well-defined use regularities. Horwich's use theory allots an important explanatory role to use regularities. In doing so it tacitly relies upon *generalizations* linking large-scale patterns of linguistic use with use regularities (in a manner that is consistent with various general psychological laws). One always encounters idealization in connection with such generalization-based explanations; idealization is the *stock-in-trade* of such explanation. It is also, at least in the context of a use theory, the source of a kind of indeterminacy that cannot be traced back to the cognitive limitations of speakers — except perhaps in the trivial sense in which the need for a theory, or an explanation, is itself the result of a cognitive limitation. Consequently, insofar as the use theorist harbors even minimal explanatory ambitions *vis-a-vis* speakers' overall patterns of use of vague predicates, he is stuck with strong indeterminacy.

VI Williamson: Bivalence by Default?

The Heap Case, I have argued, shows that there is no way for the use theorist to coordinate the meaning-constituting role of use regularities with their role in the explanation of large scale patterns of linguistic use that is consistent with the Classical View. Insofar as the Heap Case involves patterns of use of a vague predicate that seem either indecisive or incoherent, it is akin to certain scenarios examined by Timothy Williamson (*Vagueness*). As with Horwich, Williamson's most immediate concern is to argue that precise meanings can supervene upon speakers' makeshift patterns of use of vague expressions — e.g., upon speakers' sometimes conflicting and indecisive dispositions to assent to/dissent from vague sentences. Thus, Williamson observes that

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 (it is charged) to suppose, incoherently, that the meanings of our words draw lines where our use of them does not. (205)

He then considers a more specific version of the complaint:

Suppose, for simplicity, that in normal perceptual conditions any competent speaker of English refuses to classify me as thin and refuses to classify me as not thin. How could the truth or falsity of "TW is thin" possibly supervene on that pattern of use? (207)

The worry here is clear, and familiar. If the epistemicist deems 'TW is thin' true or false despite speakers' refusal to classify TW as thin or as

not thin, he extends the truth conditions of 'TW is thin' beyond speakers' use of that sentence, thus (apparently) violating the supervenience of meaning upon use. Yet Williamson replies that this apparent problem will persist only so long as one sees the values (T) and (F) as 'symmetrical.' In his view,

the concepts of truth and falsity are not symmetrical. The asymmetry is visible in the fundamental principles governing them, for (F) is essentially more complex than (T), by its use of negation. The epistemic theorist can see things this way: if everything is symmetrical at the level of use, then the utterance fails to be true, and is false in virtue of that failure (if it says that something is the case). In that sense, truth is primary, there is no symmetry to break. (*Vagueness*, 208)

Following Williamson, let's call this criterion for determining the truth values of vague sentences in borderline contexts 'Default-F,' and criteria of this general kind 'default principles.' In that case, Williamson's Default-F can be understood to entail that, if all speakers are disposed neither to assent to nor dissent from 'TW is thin,' then 'TW is thin' is false (in virtue of its failure to be true). More generally, Default-F bids us to assign the value (F) to vague sentences in all those contexts in which they might typically be believed to lack determinate truth value.²³

Williamson is circumspect in detailing the nature of default principles. In particular, he has said that his commitment to Default-F is at best tentative, and that neither Default-F nor Default-T, taken separately, is entailed by the epistemic theory of vagueness. He has also suggested that Default-F might be appropriate in some cases, Default-T in others, 'cases of each kind being too close to cases of the other to enable us to discriminate between them' ('Reply,' 951). Nor does Williamson endorse any particular account of how we come to know default principles. Indeed, he even entertains the possibility that we are incapable of knowing them or that, while we can know that there are some true default principles, we cannot know which ones they are. What we do know, in Williamson's view, is that default principles are not stipulative in nature.²⁴ On the

23 Williamson characterizes default principles as ranging over utterances, not sentences. For continuity I will continue to speak in terms of sentences, but we remain free to say that the vagueness of a sentence derives from the possibility of someone's using it to make a vague utterance.

24 At one point, after saying he is not sure how to account for our knowledge of default principles, Williamson notes that this poses no special problem for the epistemic theory and that it leaves 'plenty of scope for further inquiry into the nature of stipulation' ('Reply,' 951). However, the suggestion here simply seems to be that, in lieu of knowledge of default principles, we can rely upon stipulation to explore the space of possibilities.

contrary, Williamson refers to default principles as ‘supervenience conditionals’ that map the truth conditions of vague sentences onto speakers’ dispositions to deploy those sentences in borderline contexts.²⁵

The first thing to notice about Williamson’s example involving ‘TW is thin’ is that it presents us with an ‘indecisive’ rather than a ‘conflicted’ pattern of use. In this significant respect it differs from my Heap Case, and neither Default-F nor Default-T could eliminate the strong indeterminacy generated, for the use theorist, by the Heap Case. For reasons already canvassed, it is not open to the use theorist to conclude that ‘*a* best explains speakers’ use of “heap”’ or ‘*b* best explains speakers’ use of “heap”’ is true/false by default; for Default-T would be unintelligible in the context of the Heap Case and Default-F would be inconsistent with a use theory of meaning. But let us set this aside.

Even if we concede Williamson his default principles, this at most shows that Williamson’s epistemicism is consistent with the weak, yet durable, claim that the meaning of a vague predicate supervenes upon the *totality* of speakers’ dispositions to accept, to reject or neither to accept nor reject sentences containing it. This is an extremely general supervenience claim. It sometimes appears as though Williamson is reluctant to move beyond the general claim, perhaps even skeptical about the possibility of moving beyond it.²⁶ Hence one might view Williamson’s default principles as simply an outgrowth of the general claim that meanings supervene upon use in a way that is consistent with the Classical View. If this reading is correct, then default principles are innocuous. In particular, as supervenience claims they will operate at such a high level of generality that they will not generate any worrisome forms of indeterminacy.

But what, exactly, is the significance of such a *general* supervenience claim, and of Williamson’s default principles? Even if Williamson has shown that precise meanings could supervene upon the totality of speakers’ dispositions to use vague predicates, this fact is not sufficient, all by itself, to reconcile the Classical View with a use theory of meaning.

25 Williamson, ‘Reply,’ 951. To say that a default principle such as Default-F is a ‘supervenience conditional’ amounts to this: Where *S* is a vague sentence and *c* is a context borderline for *S*, Default-F says that *if* speakers neither assent to nor dissent from *S* in *c*, *then S* is false. Default-F is thus a supervenience conditional whose antecedent is a situated pattern of use and whose consequent is the value (F).

26 For example, he says, ‘The epistemic theory of vagueness makes the connection between meaning and use no harder to understand than it already is. At worst, there may be no account to be had, beyond a few salutary remarks. Meaning may supervene upon use in an unsurveyably chaotic way’ (*Vagueness*, 209).

For if what we are prepared to call a 'use theory' does not in any way move beyond the general supervenience claim (e.g., by invoking discrete, meaning-constituting use regularities), it will not be equipped to do any of the explanatory work of Horwich's theory. In that case, the 'use theory' seems to be less a theory of meaning than a proposed *constraint* upon an acceptable theory of meaning — and a relatively weak constraint at that since, taken by itself, it would not necessarily favor a use theory over, say, a truth conditional theory of meaning.

At other times Williamson does seem concerned to argue that epistemicism is consistent not merely with the supervenience of meaning upon use but with a use theory of meaning. For example, he invokes the notion of a 'linguistic practice,' a notion that seems to correspond nicely with Horwich's notion of a 'use regularity.' In Williamson's view, '[to] know what a word means is to be completely inducted into a practice that does in fact determine a meaning' (*Vagueness*, 211). Yet the *linguistic practice* that determines the meaning of a word is also presumably a relatively discrete item that plays an *explanatory role* relative to the linguistic use of an individual speaker or speech community. Moreover, it is presumably *in virtue* of playing such an explanatory role that a linguistic practice determines the meaning of a word. In that case we can raise all of the same troubling questions about linguistic practices that we raised about Horwich's use regularities. No doubt we will also arrive at the same answers, such as that there is a form of strong indeterminacy affecting our linguistic practices and, to go with it, a form of strong indeterminacy affecting the meanings of vague predicates. This leaves Williamson stuck with the very form of strong meaning indeterminacy facing Horwich. Hence it is misguided to conclude that the Classical View — or any theory of vagueness entailing the Classical View — is consistent with a use theory of meaning.

VII Conclusion

Horwich and Williamson have provided detailed and elegant analyses of the vagueness affecting predicates such as '... is a heap' and '... is red' as consisting in a 'paralysis of judgment' or a 'species of ignorance.' At no point in offering these analyses do they appear to say anything that is inconsistent with a use theory of meaning. In that case, it seems they ought to be able to analyze the apparent indeterminacy affecting '... best explains speakers' use of "heap"' and '... constitutes the meaning of *F*' along the same lines.

What the Heap Case shows is that these two projects are subject to radically different constraints. As mentioned earlier, the most obvious obstacle to embracing both a use theory of meaning and the Classical

View of vague language is the fact that no use regularity governing speakers' use of (e.g.) 'heap' will explain why 'heap' should have any precise extension. If the use theorist accepts a deflationary account of truth, then he can easily avoid this obstacle by holding that the use regularity governing 'heap' must *determine* a precise extension for it, but need yield no *explanation* why 'heap' has precisely the extension that it does.²⁷ Whether or not one believes that there is such an 'explanatory gap' between use and meaning, such a gap is perfectly consistent with both the supervenience of meaning upon use and a use theory of meaning.

Yet if we attempt to make a plausible parallel claim — to describe some comparable explanatory gap, paralysis of judgment or epistemic failing — in connection with the indeterminacy affecting use regularities, our effort proves pointless. If the use theorist is attracted to deflationism or is simply daunted by the prospect of deciding between various non-classical logics, he can of course *suppose* that the meaning of 'constitutes the meaning of *F*' determines some precise extension. But when he harnesses this supposition to a use theory of meaning, he commits himself to the view that the extension of 'constitutes the meaning of *F*' consists of a single, explanatorily basic use regularity (of which the predicate is true), and a host of other use regularities (of which the predicate is false). Yet, as the Heap Case shows, we have *ample* reason to deny this consequence. There is no reason to believe, and there are overwhelmingly persuasive reasons to deny, that speakers' use of *F* must be governed by some unique, explanatorily basic use regularity. In the end, there is a form of vagueness affecting the inner workings of the use theory that cannot be explained by the accounts of vagueness offered by Horwich and Williamson.²⁸

Received: June 2002

27 In Horwich's terminology, use must determine, but need not DETERMINE, meaning.

28 I would like to thank Brad Cohen, Chris Hill, Joe Moore, and especially Timothy Williamson for helpful discussion and/or comments on previous drafts of this paper. Thanks also to the National Endowment for the Humanities for providing some financial support.

