

VAGUENESS AND THE SORITES PARADOX

Kirk Ludwig & Greg Ray
University of Florida

For what you're admitting is that you cannot answer which is the last of "few" or the first of "many". Carneades

I

In "The Philosopher Is In" ("Shouts and Murmurs", *The New Yorker*, June 1998), Lawrence Douglas and Alexander George parody a philosophical counseling session in which the following exchange takes place between patient and counselor.

"Where psychoanalysis has failed, syllogism is sure to succeed. Tell me more about what's been troubling you."

"Well, there's my job."

"Yes?"

"I'm an I.R.S. auditor."

"Ahh. And what would you most like to be?"

"I've always wanted to be an orthodontist—nothing beats orthodontia."

"Let's reflect on this. You'll agree that auditing is better than nothing."

"That's certainly true."

"And you have just granted that nothing is better than orthodontia."

"Yes."

"It follows, therefore, that auditing is better than orthodontia."

"That makes me feel a little better. I'm starting to see the value of this therapy."

"Indeed, at five hundred dollars a session, it's a bargain."

"Are you nuts?"

"It's really a negligible sum."

"Not to an I.R.S. auditor."

"If I charged you merely one dollar, you'd agree that that was a negligible amount."

"Yes, of course, but—"

"And if you were to take a negligible amount and add a single dollar, you'd be left with a negligible sum, wouldn't you?"

"Well, yes, I suppose so."

“It follows, *pari passu* and *mutatis mutandis*, that five hundred dollars is a negligible sum. As Marx said, ‘Money is the absolutely inalienable—’

“I can see I’m going to have to pull out last year’s returns and—”

“Commodity. Because it’s all other commodities divested of their shape, the product of their universal alienation.’ What creates unhappiness, you see, is not unresolved childhood trauma but the absence of philosophical examination. And now I’m afraid our time is up.”

Part of the pleasure of the piece lies in the recognition that the reasoning involved, though superficially plausible, is sophistical. The patient is being bamboozled by verbal slight of hand. This is as apparent to the non-philosophical reader as to the philosophical reader, though the latter will take special pleasure in the recognition of familiar faces. The first bit of reasoning founders on the confusion of a quantifier with a noun. It is readily recognizable as a bit of fallacious reasoning, and well understood. The bit of sophistical reasoning that we are particularly concerned with in this paper is the next, the sorites argument, so-called after the ancient problem of the heap (*soros*). This is generally thought to be fallacious as well. Beyond that, however, and in contrast to the first case, there is not much consensus on what has gone wrong. This is a sure sign of a knot in our thinking whose unraveling promises fresh insight into the workings of language. Clarifying our thinking about this can also help when we are faced, as sometimes happens, with the serious use of such arguments in, e.g., ontology, the philosophy of mind, and ethics.

A sorites argument is a symptom of the vagueness of the predicate with which it is constructed. A vague predicate admits of at least one dimension of variation (and typically more than one) in its intended range along which we are at a loss when to say the predicate ceases to apply, though we start out confident that it does. It is this feature of them that the sorites arguments exploit. Exactly how is part of the subject of this paper.

The majority of philosophers writing on vagueness take it to be a kind of semantic phenomenon. If we are right, they are correct in this assumption, which is surely the default position, but they have not so far provided a satisfactory account of the implications of this or a satisfactory diagnosis of the sorites arguments. Other philosophers have urged more exotic responses, which range from the view that the fault lies not in our language, but in the world, which they propose to be populated with vague objects which our semantics precisely reflects,¹ to the view that the world and language are both perfectly in order, but that the fault lies with our knowledge of the properties of the words we use (epistemicism)²: it’s all a matter of who’s master; on this view our words have the upper hand.

In contrast to the exotica to which some philosophers have found themselves driven in an attempt to respond to the sorites puzzles, we undertake a defense of the *commonsense* view that vague terms are *semantically* vague. Our strategy is to take a fresh look at the phenomenon of vagueness. Rather than

attempting to adjudicate between different extant theories, we begin with certain pre-theoretic intuitions about vague terms, and a default position on classical logic. The aim is to see whether (i) a natural story can be told which will explain the vagueness phenomenon and the puzzling nature of soritical arguments, and, in the course of this, to see whether (ii) there arises any compelling pressure to give up the natural stance. We conclude that there is a simple and natural story to be told, and we tell it, and that there is no good reason to abandon our intuitively compelling starting point. The importance of the strategy lies in its dialectical structure. Not all positions on vagueness are on a par. Some are so incredible that even their defenders think of them as positions of last resort, positions to which we must be driven by despair of any alternative. We aim to show that there is no pressure to adopt these incredible positions, obviating the need to respond to them directly.

If we are right, semantic vagueness is neither surprising, nor threatening. It provides no reason to suppose that the logic of natural languages is not classical or to give up any independently plausible principle of bivalence. Properly understood, it provides us with a satisfying diagnosis of the sorites argumentation. It would be rash to claim to have any completely novel view about a topic so well worked as vagueness. But we believe that the subject, though ancient, still retains its power to inform and challenge us. In particular, we will argue that taking seriously the central phenomenon of predicate vagueness—the “boundarylessness”³ of vague predicates—on the commonsense assumption that vagueness is *semantic*, leads ineluctably to the view that no sentences containing vague expressions (henceforth ‘vague sentences’) are truth-evaluable.⁴ This runs counter to much of the literature on vagueness,⁵ which commonly assumes that, though some applications of vague predicates to objects fail to be truth-evaluable, in clear positive and negative cases vague sentences are unproblematically true or false. It is clarity on this, and related points, that removes the puzzles associated with vagueness, and helps us to a satisfying diagnosis of why the sorites arguments both seem compelling and yet so obviously a bit of trickery. We give a proof that semantically vague predicates neither apply nor fail-to-apply to anything. It follows that it is a mistake to diagnose sorites arguments, as is commonly done, by attempting to locate in them a false premise. Sorites arguments are not sound, but not unsound either. We offer an explanation of their appeal, and defend our position against a variety of worries that might arise about it.

The plan of the paper is as follows. In section II, we first introduce an important distinction in terms of which we characterize what has gone wrong with vague predicates. We characterize the natural starting point in thinking about the phenomenon of vagueness in section III, from which only a powerful argument should move us, and then trace out the consequences of accepting this starting point in section IV. In section V, we consider the charge that among the consequences of semantic vagueness are that we must give up classical logic and the principle of bivalence, which has figured prominently in arguments for

epistemicism. We argue there are no such consequences of our view: neither the view that the logic of natural languages is classical, nor any plausible principle of bivalence, need be given up. Next, in section VI, we offer a diagnosis of what has gone wrong in sorites arguments on the basis of our account. In section VII, we then present a simple, novel argument to show that our account must be accepted on pain of embracing (in one way or another) the epistemic view of “vagueness”, i.e., of denying that there are any semantically vague terms at all. In section VIII, we discuss some worries that may arise about the intelligibility of our linguistic practices if our account is correct. We argue none of these worries should force us from our intuitive starting point. Section IX is a short summary and conclusion.

II

We begin with an important three-fold distinction forced on us independently of considerations of vagueness. We distinguish between something’s, relative to a language *L*, *applying*, *failing-to-apply*, and *not applying* to a thing (or sequence of things—henceforth we will subsume sequences of things under ‘things’). These are semantic relations. The distinction between the latter two is of particular importance. The predicate ‘has -exactly-three-sides’ *applies* in English to all triangles, and *fails-to-apply* in English to any other sort of closed polygon. *Failing-to-apply*, however, must be distinguished from not applying. For example, a horse neither applies nor *fails-to-apply* to any polygon, in any language. Horses are not in the application business. *Applies* and *fails-to-apply* are both semantic notions, and themselves relate only things which have the requisite semantic properties, relative to a language, to other things. Horses lack the requisite semantic properties. So, while a horse *does not* apply to anything (in any language), neither does it fail-to-apply to anything.

This point extends to some linguistic expressions (understood relative to a language) as well, though we might be apt to overlook this because we are used to thinking of linguistic expressions as being candidates for having the right properties. That is, some linguistic expressions (taken relative to a language) are like horses in the relevant respects. This is unexceptionable. For example, ‘is brillig’—a meaningless predicate in English—neither applies nor fails-to-apply to anything (in English). See table 1 for illustration. Alternatively, consider the predicate letters of the artificial languages we introduce in logic classes, e.g. the ‘F’ in ‘ $(x)Fx$ ’. Such an artificial language is, by design, only semi-interpreted, so while ‘F’ is treated as filling the role of a predicate of it, ‘F’ is not otherwise meaningful. The notion of *model* that we define for these languages shows what sort of semantic value a meaningful predicate of the language is supposed to have—namely, an extension-fixing one—but the predicates of such artificial languages do not have determinate extensions. Thus, again, ‘F’ in our artificial language neither applies nor fails-to-apply to anything. The same may be said for words that are not simply nonsense words. In

<i>x applies in L to y</i>	
‘has exactly three sides’ applies in English to all triangles	
<i>x does not apply in L to y</i>	
<i>x fails-to-apply in L to y</i>	<i>x neither applies nor fails-to-apply in L to y</i>
‘has exactly three sides’ fails-to-apply in English to squares	Monarchos (the 2001 Kentucky Derby winner) neither applies nor fails-to-apply in any language to anything
	‘is brillig’ neither applies nor fails-to-apply in English to anything

Table I

‘I did it for his sake’, the word ‘sake’ neither applies nor fails-to-apply to anything in English. It is syncategorematic rather than predicative.

It is clear that we can manufacture cases of predicates which are not simply meaningless in the languages in which we introduce them, yet still neither apply nor fail-to-apply to some things. Let ‘vorpal’ be a predicate that applies (in an extension of English) to integers greater than nine, and fails-to-apply to sets of fewer than seven elements; about other things no direction is given. Whatever else may be true about ‘vorpal’, it is clear that (i) it is not simply meaningless, and (ii) it neither applies nor fails-to-apply (in our language) to Julius Caesar or the blades that killed him. In addition, arguably, many predicates have restricted ranges of application, and so do not apply or fail-to-apply to some things, though this is a perfectly determinate part of our practice using them. The predicate ‘is divisible by 2’ either applies or fails-to-apply to every number, but neither applies nor fails-to-apply to any galaxy.

Thus, it is true of each and every predicate (and everything else for that matter) that it either applies or does not apply to any given thing (relative to a language). Yet, some things neither apply nor fail-to-apply to this or that, or maybe all things (in this, that, or any language). Non-predicates in a language neither apply nor fail-to-apply to anything. More importantly, some predicates in a language may neither apply nor fail-to-apply to some given thing. Correspondingly, we require for sentences a three-fold distinction between being true, being false, and being neither true nor false, which is equivalent to being not true but not false.⁶

It is a defect of our practice in using a predicate in a natural language that it fails to determine that the predicate applies, fails-to-apply, or neither applies nor fails-to-apply to a given thing.⁷ In this case, the rules the practice determines fail of completeness.⁸ We will say that predicates in a language L are *semantically complete* provided that, relative to L, (i) a range of application is

determined, and (ii) for anything which may fall in the range of application it is determined whether the predicate either applies or fails-to-apply to it. We will say that such predicates have an *extension* in L. Other predicates in L we call *semantically incomplete*, and we say that they *lack an extension* in L.

III

Now to vague predicates. Here we mean the class of terms such as ‘bald’, ‘heap’, and so on, which are generally called vague predicates (we focus on what is generally thought of as degree vagueness), and whose characteristic is that they give rise to sorites puzzles. In this section, we aim to make a claim about such predicates (not a surprising one, we hope). Vague predicates are intuitively semantically incomplete predicates (relative to their languages). That is, *prima facie* there are some (possible) things at least, in their intended range, to which vague predicates in a language L neither apply nor fail-to-apply. Moreover, they are semantically incomplete in a rather special way. Along one or more dimensions of variation in the intended range of application, there is *no precise borderline* between what we will call fully acceptable uses of a vague predicate or its negation, and those which call forth some hesitancy. There is no saying precisely when, by successive removals of a grain of sand at a time, something we were inclined to call ‘a heap’ should no longer be so denominated.⁹ Terms like this we will call semantically vague predicates. (Our manufactured case of ‘vorpal’ is a semantically incomplete predicate that is not semantically vague. ‘Vorpal’ is an example of what we might call a *merely* semantically incomplete predicate. The meaningless ‘brillig’ is a sort of trivial example of this same category of expression.) Our unsurprising thesis (though of course it has been denied) is that vague predicates are *semantically* vague. This is the traditional, and, surely, the default position, though clarified a bit now in terms of the distinction we have drawn between not applying and failing-to-apply (as opposed to applying and not applying). Failure to mark this distinction has vitiated much discussion of vague predicates.

We think this commonsense judgment about vague predicates is correct—even obviously correct. It would take a powerful argument to move us from this starting point. Schooled in the scholarly literature on vagueness, one can lose sight of the obviousness of this observation as a starting point in thinking about the phenomenon of vagueness. No one can seriously suppose that the term ‘love seat’, e.g., was introduced into English with a precise and exhaustive set of rules determining its use with respect to any possible object one might be presented with in experience. Similarly, no one can think that someone anticipated compact disc technology when introducing ‘record album’, and it seems likewise implausible that something about early applicative (meaning determining¹⁰) practices with the term (used uniformly to apply to analog, vinyl LPs) determined back then that CDs were to be included. And so on, for the entire catalog of vague terms (most natural language predicates intended to apply to spatio-temporal objects).

Vague predicates have partially, but incompletely specified meanings. There are three factors that may come into play in the determination of a meaning for a given predicate: (i) our linguistic intentions with respect to its use, (ii) our practices in using it, and (iii) (perhaps) natural divisions to which the word is somehow keyed. All three of these factors together do not fully determine a meaning for a vague predicate, and so likewise do not determine an extension for it. We introduce expressions into the category of *predicate*, and employ them in that grammatical role. Yet, we often do not do any of the things that would be sufficient to make such expressions semantically complete. We only partially determine a meaning for these terms. Decisions that would settle the location of the breakpoints are, in Arne Naess's apt phrase, below our level of definiteness of intention (Naess 1953). It is not surprising or unnatural that this should be so (especially with "empirical terms"). It is just what one would expect given our ignorance about the world and the exigencies of life. We serve our purposes with these vague terms by dint of a subtle interplay of simple language concepts and partially specified meanings for terms.¹¹ This calls for no special explanation. It is clearly *possible* for there to be terms in a language which have this sort of semantic defect. Given the facts about the origins of words like 'tall' or 'swarthy' and how we come by the use of them, it is hard to imagine any good grounds for denying that terms with this sort of semantic defect are also *actual*. It would be astonishing, in fact, if there were not semantically vague terms in natural languages, wild shrubs that they are, untrained to the requirements of a logically ideal language.

IV

Accepting that vague predicates are semantically vague has important implications for the evaluation of sentences in which they appear. Our concept of language accords a simple semantic role for predicates—just one way of contributing to the truth conditions of sentences containing them. It is that semantic role that we describe (e.g., for monadic subject-predicate sentences) by writing (somewhat carelessly, as it turns out) something like (SR).

(SR) A subject-predicate sentence $\lceil \zeta \alpha \rceil$ is true in L iff ζ applies in L to the referent in L of α .

(We ignore context sensitivity for convenience.) With this sort of clause we describe the kind of semantic contribution that a predicate can make to such a sentence, and there are no others.

To declare an expression a predicate assigns it a grammatical role, and also determines what *sort* of semantic contribution (if any) that expression makes to sentences involving it. Yet, to do this much does not yet ensure that the expression in question has the requisite semantic properties to actually contribute semantically to sentences containing it. For example, to declare 'foo' a predicate and use it as such would confer upon it at best a *schematic* meaning. Like the

horse, 'foo' would still lack the semantic properties for it to be said correctly either to apply or to fail-to-apply to anything, and no sentence containing it would be evaluable as true or false.

To contribute semantically in the way conceptually prescribed for predicates requires that a predicate be *semantically complete*, i.e., that the predicate have an extension. Thus, the truth-evaluability of a sentence depends on the sentence being grammatically well formed from semantically complete components. We will have more to say about this claim later. Coupling this claim with our earlier claim about the semantic incompleteness of vague predicates gives us immediately the result that no sentence containing a vague term is truth-evaluable. Such component predicates fail to have the requisite semantic features presupposed by our concept of *true sentence*.¹² A predicate's lacking an extension is in this respect parallel to a singular term's lacking a referent: it fails to have a property required for it to contribute in the way required for a sentence in which it is used to have a truth-value.

This shows that our rules for the contributions of predicates to the truth conditions of sentences must be conditionalized on the predicates and other terms in the sentence meeting the conditions required for them to be truth evaluable. This is of course an obvious point in general, since when we are considering natural languages we must guard ourselves against the attempt to give truth conditions for sentences containing nonsense words. Semanticists routinely ignore this because it is useful to pretend that natural languages do not contain any semantically defective terms. Properly speaking, though, a clause like (SR) should be rewritten as (SR*):

(SR*) For any fully meaningful predicate ζ , and name α , the subject-predicate sentence ' $\zeta\alpha$ ' is true in L iff ζ applies in L to the referent in L of α .

If this is correct, then no sentence with a semantically incomplete term is true or false. This will be resisted, particularly in the case of vague predicates. It will be supposed that there is a middle position available according to which at least *some* sentences with vague predicates could have a truth-value, since the predicate does have at least a *partially* specified meaning. If our discussion in this section is correct, this is not an available position. However, in the special case of vague predicates, we can offer a proof that the middle position is not available, apart from the general considerations of this section. *The proof relies only on accepting that the boundarylessness of vague predicates is a semantic phenomenon, and the assumptions of the middle position itself.* We give this proof in section VII.

V

We accept that there are vague predicates in natural languages, and that a consequence of this is that sentences containing them lack truth-values.¹³ Must

we as a consequence give up the view that the logic of natural languages is classical?

There are two questions to be dealt with here. The first question is whether the phenomenon of vagueness gives us reason to think that classical logic is not the right logic for natural language. The answer to this question is ‘No’; classical logic still looks like the right logic for natural languages. The second question is whether vagueness forces us to “give up” in some sense using the rules of classical logic in reasoning with sentences containing vague predicates. These two questions are independent of one another. The first question is about the semantics of natural language, and, in particular, the semantics of its logical terms. The second question is about our practice of applying classical logic. We put the second question aside for now (see section VIII).

The question whether classical logic is the right logic for natural languages is really the question whether the semantics of natural languages and their logical terms (and structures) presuppose all fully meaningful (coherent) sentences (or utterances) are either true or false, that is, whether the language will receive a semantics of the general sort offered for classical logics.

The default position about this is that the logic of natural languages is classical. Our aim is to show that the existence of semantically vague predicates gives us no reason to reject this default position.

‘Logic’ serves many masters. Philosophers and others have used this term for many different things, and philosophers often mean one thing, and then another by it.¹⁴ What we seek to isolate is that sense of the term ‘logic’ which would make it sensible to argue about whether “classical logic is the correct logic for natural language.” One family of very broad uses of ‘logic’ simply equates a logic with a formal system. This is the sense of ‘logic’ in phrases like ‘logic of conversation’, ‘logic of questions’, and ‘logic of nonsense’. There is systematicity in many things, and (to a first approximation) where there is systematicity a formal system can be constructed to track that systematicity. Many things can even be tracked syntactically. In this sense of ‘logic’, there can be no question about some logic or other being *the* correct logic for natural language. There will be many logics tracking various things, some of great interest to us. Supervaluation theory is an example—it offers a structure which models (under an idealization) certain interesting facts about languages with vague expressions in them—but supervaluation theory does not give us a picture of the semantics of natural language, and, thus, it will turn out that a supervaluationist *system* does not give the *logic* for natural language.

In that sense of *logic* pertinent to our question, the logic for a language is determined by (i) the meanings of certain expressions in the language (the logical terms), together with (ii) the semantic types of the other expressions of the language, and (iii) the grammatical modes of combination of expressions. In this sense of *logic*—keyed to the semantics of the logical terms of the language—it makes sense to ask whether the logic of a language is classical or not. Clearly, there are many formal systems that track various features of a language but do not represent the logic of the language in the above sense,

because they are not concerned specifically with the semantics of the logical terms. In particular, the logic of a language in the present sense is determined by semantic features of a language which are independent of the presence of vague terms in the language.

It is of the utmost importance not to confuse modeling some aspects of a language or of language use, which involves keeping track of various features of molecular sentences on the basis of features of their included atomic sentences, with providing a semantics for sentential connectives, in the sense of providing an account of their meaning. The term 'semantics' displays some of the same elasticity as the term 'logic', and is often used in the sense of 'formal model', and this may help to explain why the literature on vagueness routinely confuses the question whether the logic of natural languages is classical with the question whether a certain model of features of natural language sentences keeps track of two or more features of them.¹⁵ Writers think of this as giving a logic of the phenomenon in question, in the sense of a formal model. They think of this as a semantics (in a sense synonymous with 'formal model'). They then think of this as giving the meaning of the logical terms of the language by an equivocation.

Consider a formal system that is designed to track various features of a language with nonsense terms in it. One thing we may wish to track is when sentences in the language are true, false, or not truth-evaluable because of the inclusion of nonsense terms. This in turn will aid us in tracking when arguments in the language lead us to true conclusions upon various assumptions about the semantic status of their premises, and whether they are valid, invalid, or simply not valid. For this purpose we may find it convenient to introduce letters, 'T', 'F', 'N', to represent something about the semantic status of the sentences of the language. This should not fool us into thinking that these letters all represent something of the same sort about the sentences we so label, and, in particular, that they represent three different semantic values the sentences may have. 'T' represents truth, 'F' falsity. 'N' represents something quite different, namely, non-truth-evaluability due to the inclusion of a nonsense term. 'The mome raths outrabe' does not have a semantic evaluation different from, but along the same lines as, being evaluated as true or false.¹⁶ To designate it 'N' is to say that it fails to come up to the standards required to be so evaluated. This is obviously not another semantic evaluation we have discovered along side the familiar. To suppose so would be to make as bad a mistake as confusing a quantifier like 'nothing' for a noun like 'orthodontia'. Thus, it is clear that a table assigning various letters among 'T', 'F', and 'N' to molecular sentences is not a representation of the semantics of the connectives appearing in those sentences, but is rather a formal representation of a status to be conferred on the whole sentence on the basis of a diverse set of features attaching to the contained sentences. That we can so keep track of features of the whole on the basis of features of the parts in this way clearly shows nothing about the meanings of the logical terms in the language, and, in particular, it doesn't show that their use does not presuppose, within the proper domain of their applica-

tion, that sentences are assigned one or another of two semantic values, truth or falsity, or that there is anything wrong with the usual characterizations of the contributions of logical terms such as the traditionally labeled truth functional connectives to the truth conditions of sentences in which they appear.

An easy way to see the point is to notice that if we were to introduce a nonsense predicate into a speech community which had by stipulation a classical logic, we would not thereby change the meanings of any of the terms already in the language or the presuppositions of their use. Knowing its grammatical category, speakers could form grammatical sentences using it. Those sentences would fail to have a truth-value. But none of the presuppositions of the use of the logical terms in the language, or their meanings, would thereby be changed. This would make the language one for which a “logic of nonsense” might be used to characterize certain features of it, but it would certainly not change the *logic* of the language, which is by hypothesis classical. The same point carries over directly to a consideration of a logic designed to track features of a language containing vague predicates (perhaps under some idealizations), such as a supervaluational theory, or a “fuzzy” logic. That we can use multi-valued “logics” (formal systems) to track features of our uses of sentences does not thereby show that the logic of the language whose features we track is multi-valued (where we have in mind semantic values on a par with truth and falsity). Just tracking some feature of a sentence that is due to a semantic defect in it is not offering a semantics for any feature of the language. Thus, it is clear that *semantically vague* predicates pose no threat to the view that the logic of natural languages is classical.¹⁷

What about the principle of bivalence? We admit sentences that have no truth-value in natural languages, so haven't we already surrendered bivalence, and with it classical logic? The answer is ‘No’. To see why, we need to clarify what the issue is. There are two separable issues here. One is about upholding the Principle of Bivalence, and the other is about admitting more than two truth values, and these may not come to just the same thing. Attend to the principle of bivalence first. What is this Principle? It is sometimes stated as follows (for a given language L):

Every (declarative) sentence of L is either true or false.

This formulation is incautious. It is unlikely that anyone ever seriously held such a principle for natural languages. First, of course, natural languages contain context sensitive terms, and so many sentences will not be thought to be truth evaluable except as relativized to a context of utterance. Secondly, even waiving this point, simple sentences with nonsense words like ‘brillig’ are quite obviously counterexamples to it. The incautious formulation above could only seem acceptable if we had already tacitly agreed to attend only to those declarative sentences that express a complete sense, or if we had already tacitly subscribed to an artificially narrow notion of the language by not counting certain terms (the defective ones) as genuine terms of the language at all.¹⁸ So, a proper

formulation of the Principle, for a context insensitive language *L*, should read something like (BV).¹⁹

(BV) Every declarative sentence that expresses a complete sense in *L* is either true or false.

For a context sensitive language, such as a natural language, we need to either relativize the truth predicate to contextual parameters, or to conditionalize on an utterance of a sentence, as in (BVN).

(BVN) Every utterance of a sentence in *L* expressing a complete sense is either true or false.

But, significantly, sentences that include semantically vague predicates (or any other sort of semantically incomplete term, for that matter) do not express a complete sense, even in use (the problem is not the need for relativization to a context of utterance). Thus, to deny truth evaluability to sentences with vague terms is not incompatible with the Principle of Bivalence, for context insensitive (BV) or context sensitive (BVN) languages.

We said that there was a second issue to be addressed when discussing bivalence. This is whether there are just two truth-values that the sentences of a language can have. It might be thought that upholding the Principle of Bivalence would already settle this question. For this to be so, (BV) would have to entail (TV).

(TV) Every sentence that has a truth-value in *L* is either true or false.

But it doesn't entail this. One needs the auxiliary assumption (X).

(X) Every sentence that has a truth-value expresses a complete sense.

While we think this assumption true, it may also appear to be question begging. Those who have proposed that there are additional truth-values have wanted to associate them precisely with sentences that have semantically incomplete terms and, hence, with sentences that don't generally express complete senses. So, while getting (TV) from (BV) requires (X), (X) appears to presuppose what (TV) was to establish. For this reason we think upholding (BV) does not settle the two-values question. Obviously, this question is run together with the one about the bivalence principle.

Regarding the two-values question, we count ourselves among those who find the assertion that there are truth-values in addition to truth and falsity unintelligible (see (Haack 1980) for defense of this position). We count (TV) a conceptual truth (for given *L*). It should be clear from things we said earlier that we don't think this point contravened by the fact that one can construct "multi-valued" formal systems that track interesting features of language use.

And, as we have said, we propose a standard classical logic for natural language as the right logic. That logic is founded on just two truth-values, and is consistent with the Principle of Bivalence, (BV).

Lest this be thought to be question begging, recall the dialectical stance of this paper. The natural stance is that the logic of natural languages is classical. The claim that it is not requires a powerful argument. Our aim has to be to show that no such argument is forthcoming from the view that vague predicates are semantically incomplete, and that consequently sentences containing vague predicates are not truth-evaluable. (Note that even if one rejects our blanket claim about this, the view that some sentences at least which contain vague predicates are not true or false is widespread. What we have argued applies equally well here. There is no need on these views either to give up the principle of bivalence or the view that the logic of natural languages is classical.)

VI

What then of the sorites arguments? If the logic of natural languages is classical, are they not valid arguments? Can any of their premises be comfortably denied? Consider the *falakros*, the paradox of the bald man.

- [1] A man with 0 hairs on his head is bald.
- [2] If a man with n hairs is bald, then a man with $n + 1$ hairs is bald.
- [3] Therefore, for any n , a man with n hairs on his head is bald.

Is the argument form invalid? If not, since we reject the conclusion, *which premise do we say is false?*

No one should be forced to answer *this* question. It has a false presupposition. Sorites arguments with vague predicates are certainly not sound. If they were, then a contradiction would be true, which is impossible. But that does not mean that they are *unsound*. For an argument to be unsound, it must be invalid or have at least one false premise. The trouble with the sorites arguments is that although they have a valid argument form (in the way that an argument form using schematic letters for predicates can be said to have a valid argument form), none of the lines of the soritical derivation is truth-evaluable (and hence none are false).

This is enough to ensure that sorites arguments are without force, but there is more that needs to be said in order to explain what is going on with sorites puzzles, since on our view many natural language arguments will fail to be sound due to the inclusion of vague terms, but few of these will be paradoxical or puzzling. Something else is going on that gives the sorites paradox its particular bite.

There is no puzzle about why we find the minor premise acceptable and reject the conclusion. The acceptance of the minor premise and the rejection of the conclusion are clearly a feature of our usual practice, and are *deliberately* chosen so that they are. The puzzling part of sorites arguments, and what clearly

generates the problem, is the step premise. We have an account of why we are inclined to assert the minor premise, and why we are inclined to assert the denial of the conclusion. What is the account of the inclination to accept the step premise, since it is so clearly not a part of the standard practice?

Our inclination to accept the step premise finds its ground in our (unarticulated) recognition that 'bald' is semantically vague, that is, in our recognition that our practices in its use give us no guidance with respect to such small differences along the relevant dimension of variation. We mistake the step premise for an expression of this fact. That is why it seems compelling, for it seems to express an intuitive truth about the use of the term 'bald'.

This mistake comes about through the failure to distinguish two different, easily confused claims about our practice in using the term 'bald'. We learn the use of 'bald' with respect to examples that are easily discriminable, and differ quite significantly, nearly hairless and hirsute heads. Our practices, however, are silent about changes on the order of one or two hairs. Significantly, it is not that our practices *tell us that* small changes make *no* difference. That would suggest that there were some rule of usage to the effect that if someone has n hairs on his head, a few more or less makes no difference to whether 'bald' applies to him. But our practices sanction no rule to that effect and are in fact incompatible with it. We are trained only with respect to examples that differ significantly; the training results in no rules that attach to particular numbers of hairs. The question doesn't come up. If it does, no one knows what to say. This explains why vague predicates are boundaryless. But it is easy, particularly in the context of a sorites argument, to mistake this silence on whether small changes make a difference for a positive rule to the effect that if 'bald' applies or fails-to-apply to a man with n hairs on his head it applies or fails-to-apply to a man with $n + 1$ or $n - 1$ hairs on his head. This is to confuse (A) with (B).

- (A) Our practices in the use of 'bald' do not give us guidance about the application of the term with respect to incremental differences in the number of hairs on someone's head.
- (B) Our practices tell us that if 'bald' applies (or fails-to-apply) to someone with n hairs on his head, it applies (or fails-to-apply) to someone with a few more or less.

(A) is true. (B) is not. The importance of keeping these distinct cannot be over-emphasized.²⁰ The step premise seems compelling because we recognize that (A) is true. Then we confuse it with (B). This is a natural enough thing to do. We are encouraged by the fact that when we feel comfortable in applying 'bald' to someone, we know that a few hairs more or less won't decrease our level of comfort (it would make no easily discriminable difference). We imagine this to be an expression of positive guidance by our practices with respect to small numbers of hairs. But it obviously is not. We are guided, to the extent that we

are, by gross appearance, and then only roughly with respect to cases near our paradigms. Once we have misrepresented the facts about the practice as in (B), it seems acceptable then to express it more precisely as:

- (C) 'bald' applies (fails-to-apply) to someone with n hairs iff 'bald' applies (fails-to-apply) to someone with $n + 1$ hairs,

Then we employ semantic descent incautiously and unreflectively to arrive at a material mode claim,

- (D) Someone with n hairs is bald iff someone with $n + 1$ hairs is bald,

one half of which is [2] above, and the other half of which can be employed in the argument that is the mirror image of the *falakros*.

We are helped to this, of course, by the fact that we are caught off guard. The step premise is introduced after a premise in which 'bald' is used in the material mode with respect to a clear positive case. We are invited to assume our usual pretense that the term is semantically complete, and that there is no difficulty in carrying on in the material mode. (This encourages us to put aside our implicit knowledge that our practices are silent about some aspects of the applicability conditions for the application of 'bald'.) Then we are presented with a premise in the material mode that seems to express something true about our use of the predicate. Indeed, continuing with the assumption we have been invited to accept as part of the usual pretense that the term is semantically complete, it will be difficult to see any way to avoid the confusion of (B) with (A), for (A) is incompatible with the assumption of semantic completeness. The difficulty in seeing clearly what is going on is compounded by the usual way of trying to enforce acceptance of the step premise: it is pointed out to us that *denying* the step premise would force one to say that there was some breakpoint number of hairs, and we will be asked what it is. Since on reflection we realize that there being a breakpoint is not in accord with our honest beliefs about 'bald', we feel compelled to accept the step premise. In short, we are sadly tricked.²¹

The confusion is a natural one. It is a confusion, because the step premise does not express the vagueness of 'bald'. It fails to express anything truth-evaluable. The truth expressed in (A) motivates its acceptance through failure to distinguish that truth from the falsehood expressed by (B). So, our acceptance of the step premise is the product of our intuitive recognition of the vagueness of the predicate, and failure to distinguish between no guidance and negative guidance. (B) is reformulated as (C). Notice that, if we are right, (C) is not in fact false, since given that 'bald' is vague, each side of the biconditional is false. So curiously we can see the implicit reasoning supporting the step premise going through a true premise. But semantic descent is illegitimate, precisely because 'bald' is semantically incomplete. Thus, with an unfortunate slip be-

tween the formal and material modes of expression, we move from a truth to a non-truth, and end up with misbegotten expression of a dimly recognized truth. The slip is natural, since (i) we have a practice of applying the predicate in material mode, and (ii) the semantics of semantically complete predicates (under which pretense we use vague predicates) warrants moving freely between the two, and, moreover, of course, (iii) (B), the motivation for (C), carries with it the (false) presupposition needed for semantic descent.²²

Sorites arguments do not just present us with a paradox. One has the strong sense when being walked through a soritical argument that it is a “set up”, a trick of some kind, even if it is not apparent how one is being tricked (we feel there is something going on we have failed to notice). In this way, it is different than, say, the Liar. Our analysis of the sorites makes this suspicion entirely explicable. The acceptance of the minor premise and the denial of the conclusion both enjoin us to engage in a certain common practice of pretense. Yet acceptance of the major premise is motivated by recognition of a fact that requires us to drop that pretense. Recognizing that fact and trying to maintain the pretense is what makes the major premise seem difficult to reject, and leads to the sense of paradox. To place both kinds of premises in an argument is a kind of trickery, albeit instructive.

VII

The diagnosis of the sorites reasoning we have just offered is grounded on the claims that (i) vague predicates are *semantically* vague, and, hence, (ii) sentences in which they appear are not semantically evaluable as true or false. This runs counter to many discussions of vagueness in the sense that often philosophers will assume that it is unproblematic that vague predicates do apply in some cases and fail-to-apply in others. This is what we earlier called the middle position. But it is not a tenable position. One *must* either accept our position, deny second-order vagueness (and explain away our inability to say where the second-order borderline is by appeal to “epistemic vagueness”), or deny first-order vagueness (and explain away our inability to say where the first-order borderline is by appeal to “epistemic vagueness”). The latter two positions are implausible on the face of it, and, if we are right, there is no pressure to accept either by reflection on the consequences of semantic vagueness.

Our argument has proceeded on the basis of general considerations about the presuppositions of our use of semantic terms, specifically, about the proper range for their application. We said earlier, however, that in the special case of vague terms, the conclusion admits of proof on the basis of the boundarylessness of vague predicates, and assumptions that the middle position theorist accepts. We present and defend the proof in this section.

The argument rests on the following theorem (‘applies to n ’ is short for ‘applies to a man with n hairs on his head’).

THEOREM:

- IF
- (i) for all n , 'bald' applies to n iff 'bald' applies to $n + 1$, and
 - (ii) for all n , 'bald' fails-to-apply to n iff 'bald' fails-to-apply to $n + 1$,
- THEN
- (iii) 'bald' applies to every n ,
 - (iv) 'bald' fails-to-apply to every n , or
 - (v) 'bald' neither applies nor fails-to-apply to any n .

A proof of THEOREM given in the appendix. But it is not difficult to see why this is true. The conditions in the antecedent require that for any n , 'bald' applies to every n -haired head if any, and fails-to-apply to every n -haired head if any. Since 'bald' cannot both apply and fail-to-apply to the same object, we immediately have that it applies to every head no matter the number of hairs, or fails-to-apply to every, no matter the number of hairs, or neither applies nor fails-to-apply to every head, no matter the number of hairs.

It may already be apparent to the reader what the implications of the THEOREM are, but before we draw them out, we want to set aside two sorts of concerns that might arise about it. First, our argument for THEOREM uses classical, first-order logic, and, of course, many philosophers working in this area have given up classical logic as a part of dealing with the issue of vagueness. But this is no objection to our argument, in which we get to stipulate that we are using classical logic. Clearly there would be no obstacle to presenting the argument in a formal language in which we give an explicitly classical semantics for the logical terms employed. This objection then is beside the point.

Second, an inspection of the argument in the appendix for the theorem reveals that it employs a little-by-little argument, and this will seem suspicious to some, since it will appear that we are using a sorites argument to get the result. Are sorites arguments not suspect? But not every argument that has the form of a little-by-little argument is a sorites argument. A little-by-little argument that employs a vague term in pride of place is indeed soritical, and thus suspect, even paradoxical. Yet, the argument for the above result does not so employ a vague predicate, and it clearly is not paradoxical in the way a genuine sorites argument is. Soritical arguments have the *form* of a (classically) logically valid argument. Thus, a little-by-little argument that does not employ any vague words (nor words semantically defective in some other way) will be an undeniably valid argument. Our argument for this theorem is of this sort (modulo a qualification we will deal with presently). So, there should be no question about the veracity of the theorem.

Our claim that the argument is not soritical may not carry conviction. That it is not hinges on whether 'applies' and 'fails-to-apply' are themselves precise. And it may be objected that there is a good argument for thinking that the term 'applies' is itself vague. For if we have a borderline case, b , for 'bald' then it is at least as problematic to assert, "'bald' applies to b " as it is to assert, ' b is

bald', and so 'applies' must be at least as vague as 'bald'; similarly for 'fails-to-apply'. Or so the argument goes. This should immediately excite our suspicion, however. For it looks very much as if the claim is that the vagueness of 'bald' is being said to be transferred to a predicate we seek to apply to it.

This is to make a mistake akin to that of assuming that if a term is used with a vague term, then it acquires the vagueness of the term with which it is used. This is to charge guilt by association. The vagueness of a term no more rubs off on terms with which it is combined in sentences than does the meaninglessness of a nonsense predicate. It would be an obvious mistake to infer from the fact that 'The mome raths outrabe and someone stole the tarts' is not fully meaningful because 'mome raths outrabe' is nonsense to the conclusion that 'the', 'and', 'someone', 'stole', and 'tarts', were likewise nonsense. Similarly in the case of sentences containing vague terms. For example, it would be a mistake to assume that because one cannot say precisely when two rods are exactly the same in length, that our word 'length' must be vague. The vagueness of 'rod' is sufficient to explain the difficulties we encounter here (the problem of the many); we have no reason to think that 'length' is similarly vague. Likewise, inferring that 'object' or 'existent' is vague because our sortals are vague would be a mistake.

We say the mistake is akin to rather than the same as the one just mentioned because it does not involve inference from use of a term in a sentence with a vague term to the vagueness of the associated term. In the above argument, the vague term 'bald' is not used in the sentence, but mentioned. But a similar invalid inference is made. It is an inference from the hesitancy over whether to assert a sentence in which a vague word is mentioned which arises from the vagueness of the mentioned word to the conclusion that another word in the sentence applied to it is vague also. How does this arise? In everyday life we act as if our predicates are precise, that is, we engage in a practice in their use that ignores their semantic deficiencies, and this is second nature to us. When we are asked to engage in semantic ascent, we typically do so without thought to whether we should abandon our tacit practice of treating our predicates as if they were semantically complete. The fact is that we are often just not very reflective about our practices or very clear headed in thinking about them. (Think of the ubiquity of use-mention confusions in ordinary speech, and even in the writings of trained philosophers: the fault lies in us, not in our semantics. We often say things that are confused and can't be taken at face value: but we can uncover the confusion if we are careful.) When we engage unreflectively in semantic ascent and apply a semantic predicate to a vague term such as 'bald', we try to let ourselves be guided in the application of the semantic predicate to a pair \langle 'bald', x \rangle ; by the practices associated with 'bald'. When the practices do not determine an extension, and no precise borderline, we find ourselves hesitating in applying 'applies' in the same way as we do in applying 'bald'. But this is because we are laboring under the (working) assumption that the predicate is semantically complete. We try to let the practices

guide us, expecting that they will, but they fail us. This does not mean that 'applies' is vague. It means that it was a mistake to try to let the practices associated with the predicate guide us with respect to the application of 'applies' to the pair of it and an object, given that it is vague, and so semantically incomplete. Rather, we must ask ourselves, in full light of the fact that a vague predicate fails to have an extension, whether it makes sense to say that it applies (or fails-to-apply) to anything. We must give up the unthinking practice of treating them in most contexts as if they were precise. When we employ semantic ascent, and step back from our practice, we can then recognize that when a predicate ζ is semantically incomplete, there is no pair $\langle \zeta, x \rangle$ to which 'applies' applies, and likewise no pair $\langle \zeta, x \rangle$ to which 'fails-to-apply' applies.²³

Of course, we do not deny that a term such as 'applies' *could* be vague in some language. But there is no good reason to think it is in our language. The evidence used to support the claim can be accounted for without supposing our semantic predicates are vague. This neutralizes the supposed evidence. And there is clearly no difficulty in seeing how to understand our semantic predicates so that they are precise, and this corresponds to their use in the semantic tradition. We can no more be forced to accept the vagueness of our semantic terms by being forced to answer the question whether 'applies' applies or fails-to-apply to the pair of a vague predicate and an object than be convicted of animal cruelty by being forced to answer the question whether the kittens have finished baking.²⁴

We have heard this objection so often that it is worth dwelling a bit more on why it misfires. There are two points that should be emphasized. First, in the present context, an object does not get the claim that 'applies' is vague for free. It is not *prima facie* a vague predicate. 'Applies' applies to a pair just in case the second member falls in the extension of the first. 'Applies' fails-to-apply to a pair just in case the second member does not fall in the extension of the first. Otherwise, it does not apply. A term has an extension only if it is semantically complete, i.e., it has a range of application and for everything in its range either that thing is in its extension or not, i.e., it is true of that thing or false of it. There is no threat of vagueness here. Indeed, it follows from this characterization that there is a determinate answer for any pair whether 'applies' applies to it, fails-to-apply to it, or neither applies nor fails-to-apply to it. Significantly, 'applies' is not an empirical term. Empirical terms are typically vague because it is hard to anticipate all the questions that arise about their application conditions (and rather pointless to do so for ordinary purposes). The development of the atomic theory of matter, for example, posed unanticipated questions about the use of virtually all of our ordinary material object terms. It is clear that the contexts in which we learn and ordinarily use these terms do not call on us to decide exactly which molecules are to be included in the extension of 'is a part of that F', and no guidance is provided by the practice as to exactly which of the more or less appropriately related molecular constituents of medium sized objects are to be counted as in the object and

which not. We learn most ordinary empirical predicates with respect to easily discriminable cases, and application conditions with respect to differences that are not easily discriminable, or which have not been anticipated, are left open. No such difficulties intrude in the case of semantic terms, and, as we have seen, barring any independent reason to think semantic terms vague, their use in application to vague terms does not show that they are vague. We may be fooled momentarily by failing to abandon our practice of treating vague terms as if they were precise when we engage in semantic ascent when doing theoretical work. However, when we take account of the false presupposition, there is a determinate answer to the question whether ‘applies’ and ‘fails-to-apply’ apply to pairs of those terms and objects or not. They do not. Rather, they both fail-to-apply, as we have said. So the answer is not in fact indeterminate at all. Our practices answer the question quite decisively when we give up our ordinary pretense that the terms to which we apply ‘applies’ and ‘fails-to-apply’ are precise.

We have been making some claims here about our ordinary semantic concepts. We believe these claims are correct and that no good argument has been advanced against them. But it is important to note that even if we were wrong about ordinary semantic terms, this is no threat to our conclusion. In theoretical work, where we find terms that are not semantically complete and which thereby hinder our progress, we introduce terms that remedy the deficiencies. If ordinary semantic terms are to be thought vague, let us then repair this deficiency to produce terms that lack these defects. Precise semantic predicates will have sharp boundaries between applying, failing-to-apply, and neither applying nor failing-to-apply, and all pairs of vague terms and objects must perforce land in the latter category. Let us stipulate then that the semantic terms used in the above argument are precise, and so operate in the way we have indicated. Then it is a determinate fact about their usage that they fail-to-apply to any pairs of vague terms and objects. The threat to the argument by the inclusion of a vague term in a crucial place is thereby met, and the argument goes through.²⁵

Perhaps it will be objected that this is to change the subject, if what we were interested in was the question whether in the ordinary sense of ‘true’ and ‘false’, any declarative sentence containing a vague term fails to be either true or false. But this is a mistake. Precising a term is not changing the subject. It is improving the focus of the discussion. We lose only unclarity. It is a mistake to think that part of the subject of what we intend to be talking about in using a vague predicate is what the vague term leaves undetermined about its use. This would be like thinking that part of what you see when you take off your glasses is the relative increase in fuzziness of your visual experience.

There is a related objection. That is that although the argument for THEOREM is not a soritical one, it and the conclusion nonetheless contain vague predicates (not ‘applies’ it will be granted), which on our account means that its premises and conclusion are not truth evaluable. This is not the term ‘bald’, of course, which is not used but only mentioned in the argument, but rather the

terms we have used to describe the series of cases with respect to which we are considering whether ‘bald’ applies or fails-to-apply. Recall that “‘bald’ applies to n ” is an abbreviation for “‘bald’ applies to a man with n hairs on his head’. The terms ‘man’ and ‘hair’ and ‘head’ are all vague terms, and perhaps ‘on’ as well. Does this not show that the premises of our argument for the THEOREM and the THEOREM itself are not truth evaluable, and, hence, that the argument on our own showing is not sound? The answer is: yes, but it doesn’t matter. The use of vague terms in indicating the sorites series of concern is clearly incidental, and is no barrier to getting the point. All that is needed is that there are or can be sorites series for the term ‘bald’. *This* is undisputed.²⁶ We use vague terms to indicate the relevant series because we are not supplied with precise terms for the job. But the argument can be restated without the use of vague terms by the simple expedient of quantifying over members of a sorites series for ‘bald’, $b_0, b_1, b_2, \dots, b_k$, and reinterpreting “‘bald’ applies to n ” as “‘bald’ applies to b_n ”. We set this worry aside then, and continue in our blithe way to employ outside the center of attention terms that are admittedly vague. We will have more to say about this practice later.

Let us return to our theorem now. Assume we have established it as we have claimed. What can we take it to show? Suppose the antecedent is true. Then we can conclude that (iii) ‘bald’ applies to every head, no matter the number of hairs on it, or that (iv) ‘bald’ fails-to-apply to every head, no matter the numbers of hairs on it, or that (v) ‘bald’ neither applies nor fails-to-apply to any head (or anything else). (iii) and (iv) are options which we take it no one is willing to accept. Each is completely incompatible with our practices. More importantly, though, for our dialectical purposes, (iii) and (iv) are ruled out for anyone who thinks that there is a middle position, according to which (a) and (b) hold given that (c) does.

- (a) There are some cases to which ‘bald’ does not apply.
- (b) There are some to which ‘bald’ does not fail-to-apply.
- (c) There is no case to which any term both applies and fails-to-apply.

(c) is a conceptual truth, since that something fails-to-apply to a thing entails that it does not apply to it. Given (c), ‘bald’ does not apply *and* fail-to-apply to any n . Given (a), therefore, (iii) is ruled out. Given (b), (iv) is ruled out. That leaves (v) as the only viable option, which is the position we have argued for independently.

Now let us ask what reason we have to accept the antecedent, the conjunction of (i) and (ii), repeated here.

- (i) For all n , ‘bald’ applies to n iff ‘bald’ applies to $n + 1$;
- (ii) For all n , ‘bald’ fails-to-apply to n iff ‘bald’ fails-to-apply to $n + 1$.

On our position, of course, it is vacuously true, since ‘applies’ and ‘fails-to-apply’ alike do not apply to any pair \langle ‘bald’, x \rangle , and thus the biconditionals (i)

and (ii) are vacuously true. Our position is thus self-verifying. But more importantly, someone who does not want to accept that position finds himself forced to choose between it and a more unpalatable position, which requires a complete rejection of the intuitively compelling view that vagueness is a semantic phenomenon. What our argument thus shows is that if one accepts that the boundarylessness of vague predicates is a semantic phenomenon, one must accept that no sentence containing a vague predicate is truth evaluable. For to deny the truth of the conjunction of (i) and (ii), one would have to hold that there is a precise borderline between cases in which ‘bald’ applies/fails-to-apply or does not apply/fail-to-apply. That is to say, to deny (i) would come to saying that for some n , ‘bald’ applies to n but does not apply to $n + 1$; to deny (ii) would come to saying that for some n , ‘bald’ fails-to-apply to $n + 1$ but does not fail-to-apply to n . This is to deny that a term like ‘bald’ is vague at all.²⁷

That is to say, one would have to accept that one of the positions (a)–(d) in figure 1 obtained, where the vertical bars represent borderlines. (a)–(c) represent cases in which there is a second-order borderline along the relevant dimension. In (a) there are second-order borderlines between both applying and neither applying nor failing-to-apply and between neither applying nor failing-to-apply and failing-to-apply. In (b) there is a second-order borderline between applying and not applying. In (c) there is a second-order borderline between failing-to-apply and not failing-to-apply. In (d) there is a first-order borderline between applying and failing-to-apply. To deny the conjunction of (i) and (ii) is to accept that one of (a) through (d) obtains. And this is to deny the boundarylessness of vague predicates, arguably the central feature of them, and what makes them both so puzzling and interesting. Denying the central phenomenon associated with a puzzle is neither to appreciate it nor to solve it.²⁸

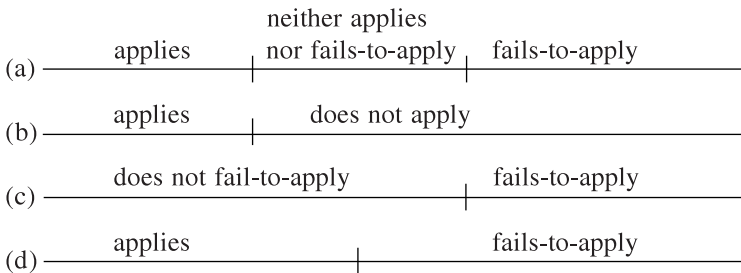


Figure 1.

This position has defenders, epistemicists, who believe that there are no semantically vague terms, and that the problem is instead epistemic. To their credit, they occupy not a halfway house, but opt boldly for (d). One might as well hang for a sheep as for a lamb. Despite valiant efforts, however, they manage only to trade a puzzle for a mystery, namely, how we could not in principle

discover boundary lines. The extensions of our predicates are fixed, when they are fixed at all, by our collective intentions and dispositions to use them, and, perhaps relative to these, some facts about the actual world.²⁹ No one doubts that there is an intelligible connection here, if there is one at all. Yet, we cannot conceive of any way in principle to settle the question at what point along a sorites series we move from someone to whom 'bald' really applies to one to whom 'bald' really does not apply. Epistemicists must then maintain that the facts on which the semantic facts supervene are inaccessible to us, or the connection is inscrutable. No epistemicist has succeeded in explaining how this could be.³⁰ Indeed, so far as we can tell, there is nothing a being omniscient about our practices (and the world) would know that would determine a cutoff point for vague predicates like 'bald'.³¹ It is not our purpose here, however, to argue directly against epistemicism. If we are right about the dialectical position, it is enough to show that epistemicism is not forced on us by finding that the natural view that vagueness is a semantic phenomenon encounters difficulties greater than those which face epistemicism.

VIII

It remains then only to assuage the anxieties that arise in some philosophers at the thought that many of our terms are genuinely semantically vague. There are various strategies that are employed to avoid these anxieties. Some deny the obvious. Some maintain the consequences of accepting semantic vagueness are not consequences of it. It is surely better to accept the obvious and its consequences, and consider whether there is anything puzzling about our practices using vague predicates.

The central worry is engendered by the thought that the sentences containing semantically vague predicates are not semantically complete. If they are not semantically complete, then we fail to assert truly or falsely using them, and, hence, are reduced to uttering nonsense, and must fail to be able to carry out any of the ordinary activities in which the use of language is so evidently useful a tool.³² But this is a confusion, a false dilemma. The options are not: assert fully meaningful sentences or assert nonsense. Vague sentences are grammatically well formed, and vague predicates have at least partially specified meanings. That is how we are able to operate effectively with them. Vague predicates have a robust practice attached to them, and it could be sharpened if there were profit in doing so. Indeed, as soon as someone points out the respects in which a term is vague, we are in a position to sharpen to within the degree allowable by the terms used in the statement of the problem. If the problem has to do with the number of hairs on someone's head, we can pick a number as the cut off point. Of course, terms we use to characterize the dimension of vagueness of a given term may themselves be vague, but then we can repeat the procedure with them in turn. Our practices in this fashion can be indefinitely sharpened, though whether it would be useful or practical to do so is

another question. Given the practice in the use of a vague term, when someone uses the term in accordance with the practice, and we know this, then, depending on how trustworthy we suppose him to be, we can learn something useful as a result, and obtain some guidance for action. We do not speak truly or falsely using vague predicates, but we are very far from uttering gibberish.

The fact that the sentences in which we use vague predicates are grammatically well-formed means that we are in a position to know what sort of truth conditions they would have (or fulfillment conditions for non-declaratives³³), were predicates in them semantically complete. This is akin to what one knows about a sentence in a formal language of the sort studied in a course in symbolic logic. Such languages are semi-interpreted—all but the predicates and names are interpreted. (This should provide one with a hint that we are also in a position to know with respect to vague sentences what sort of logical consequences they would have, if their predicates had extensions.)

So, we know what sort of truth conditions these sentences *would* have. We also know something more than this, because we have at our disposal whatever we know of what there is to know about the partially determined meaning of a vague predicate. This may amount to such things as (i) knowing what are supposed to be determining factors for the application of the predicate (quantity of hair on the head, for example), and (ii) some ranges of such determining factors which are intended cases for the predicate.

Such knowledge of our partially formed semantic intentions for the vague predicates involved, together with knowledge of the semantic type of the sentence, give us a great deal to go on in communicating with each other—more than enough to make use of these vague terms useful in just the circumstances in which we do use them.

The trick to seeing why we should not feel threatened by the pervasiveness of vague terms in natural languages is just to see how natural and sensible it is that practices using such terms would have developed between us. Our picture is that we have simple language concepts, and that we develop terms that imperfectly meet the requirements for the application of these concepts. So, we have these predicates that don't have all the semantic properties necessary for the application of the notion of truth and falsity to sentences containing them. However, by treating these words in a context as though they fit the mold, and deploying them according to a common practice that we have with them, we manage to get various things done that we want done, and various things communicated that we want communicated. The truth is that it would far exceed any ordinary purpose that we have to make our predicative terms comply with our ideal. Partly, this is because we are always simplifying the domain about which we wish to think. So, in fact, the practice that we have is well suited to this, since it allows a large family of terms to go flush with that habit of simplification. As Quine has put it, our “terms delimit the object to the degree relevant to our concerns” (Quine, 1985, p. 168). For example, it would not do to define ‘human’ so as to have precisely circumscribed molecular boundaries,

i.e., so that, in principle, each molecule in the vicinity was determinately in the extension of 'is a part of that human' (as used on a particular occasion). This is because we do not ordinarily conceive of our "humanish" surroundings in these terms, and it would moreover be impractical to do so. We conceive of there being just these few largish humanish things, complex things with parts, to be sure, but things that come in large chunks, not stacks of atoms. Generally speaking, terms like 'human' are good to use when certain kinds of simplifications suit our purposes, and not good for others. Thus, if we are interested in the hunting practices of a group of humans, then a model of our domain that included items representing humans, but no items representing molecules, might well serve our purposes. Whereas, if our interest was in the migration of carbon atoms across a certain type of membrane, then a model of our domain which included items which represented atoms and molecules would be indicated, and, as like as not, the term 'human' will have no easy or relevant application.³⁴

Another worry that arises about admitting that many of our predicates are vague is that this will have ontological consequences that no one will wish to accept. Thus, for example, Peter Unger has argued that the phenomenon of vagueness forces us to the conclusion that there are no people or, indeed, any ordinary things at all (see (Unger 1979b) and (Unger 1979a)). Unger argues that our practices commit us to the step premises and minor premises of sorites arguments, but also to the negation of their conclusions, thereby demonstrating that the rules governing vague predicates are inconsistent. From this he concludes that 'person' applies to nothing, and, more generally, that most of our ordinary terms apply to nothing. This is a conclusion of course that we are committed to independently. This only gets us that a term like 'person' applies to nothing. But the argument from this conclusion to the further conclusion that there are no people is simple.

'person' does not apply to anything
Therefore, there are no people.

We have argued that it is a mistake to think that our practices commit us to the step premise of a sorites argument. But we are committed to the premise of this argument anyway. Clearly, if this argument is sound, then so will be any argument of the same form for any vague predicate one cares to choose. If this argument is correct, then not only are there are no people, there are no chairs, no trees, no mountains, no nations, no stars, etc. This is surely an intolerable result. Fortunately, the argument is not sound. While the first premise is true, the conclusion is neither true nor false, since it contains a vague term, 'people'. In general, the argument form,

'F' does not apply to anything
Therefore, nothing is an F

is valid only if 'F' is semantically complete. If 'F' is a vague term, it is not, and so the argument schema will not yield a sound argument. Our view therefore has no ontological consequences whatsoever.

Must we refrain then from asserting sentences containing the word 'people' altogether? Of course not. When we are participating in our usual practice, it is perfectly appropriate to assert 'There are people', and inappropriate, indeed, outrageous, to assert 'There are no people'. We do this under the pretense that the terms are precise, and under that pretense, it is appropriate to assert 'There are people', and inappropriate to assert its negation. What is illegitimate is to step outside of the practice, recognize the terms are vague, and then attempt to carry on as if we had not given up the pretense that the terms are semantically complete.

This defense of the ontological neutrality of our position just highlights what is apt to be another worry, however. Surely, on our view, there are virtually no good arguments in natural languages. For no argument containing a vague predicate on our view is either sound or unsound. They are all on a par in simply being not sound. We are then surely left with the unenviable position of not being able to explain the robust distinction between good and bad natural language arguments, and unable to explain their undoubted utility. But although it is a consequence of our position that natural language arguments containing vague terms are neither sound nor unsound, it does not follow that they are all on a par, and that we cannot account for the utility of those arguments we naturally think of as good arguments.

A first observation is that though we count no argument containing a vague term sound or unsound, we can still distinguish between arguments valid in virtue of their form and arguments that are not valid in virtue of their form, for this abstracts from the content of the predicates in the arguments. So it is simply not true that we have no means of distinguishing between the merits of natural language arguments. In very many cases, we will use exactly the same criterion as anyone else.

There still remains to explain how arguments with premises and conclusions that are not true because of the inclusion of vague terms may be useful. The answer, in outline, is as follows. Although these arguments do not preserve truth, because their premises are not true, they do preserve something else, which we will call appropriate assertability. To a first approximation, if the premises of an argument are appropriately assertable, then if the argument would be valid on all acceptable precisifications of its vague terms, the conclusion is appropriately assertable.³⁵

We will come back to this in a moment. First, though, what do we mean by 'appropriate assertability', and why is this to be linked to the usefulness of assertions of sentences containing vague predicates? Securing the link is the easy part, for we wish to understand the notion of appropriate assertability in terms of assertion in conditions which make the assertion useful for the purposes of conveying information about the world to an interlocutor. The trouble is to say something informative about those conditions. We use sentences con-

taining vague predicates to convey information about the world. That information is conveyed by the fact that we sincerely assert the sentence (and so can be taken to think that our assertion of it falls well-within the practice of use of the vague predicates it contains) and the practice associated with the terms in the sentence, including the vague terms. An auditor gains information about the world by way of knowledge of our intention to use the predicates well within the practice (this will be the default case, though there will be others) and knowledge of the practice associated with the predicates. This will give rise to expectations about what is to be encountered in experience if the speaker's beliefs about the world are taken to be correct. Typically, unless someone is simply drawing our attention to something salient in the environment, the expectations engendered will be degrees of belief attaching to a range of possible states of affairs with a profile determined by the degrees of vagueness of the predicate.

Let us take a simple example. Suppose that someone asserts 'Bert is bald' of some person we have not yet encountered, but of whom perhaps we know some things. Let us for convenience continue with the simplifying assumption that amount of hair on the head is the only relevant dimension of variation for application of the term. In hearing our interlocutor assert 'Bert is bald', we are not thereby given information about how many hairs to expect Bert to have on his head, of course, nor even anything very precise about the extent of uncovered pate on his head. What we learn, on the assumption that our interlocutor intends to be using the term in accordance with our usual practice, is (roughly) that given a range of relevant cases of heads from one with no hair, b_0 , to one completely covered, b_k , the probability that Bert is similar in numbers of hairs to b_0 is high, and the probability that Bert is similar in numbers of hairs to b_k is low. This gives us useful information on which we can act as appropriate. We will say therefore that a sentence containing a vague predicate is appropriately assertable provided that the expectations that its sincere assertion, taken non-figuratively, would thereby standardly generate, on the assumption that the speaker is using the terms in it in accordance with the practice, will not be frustrated.

If appropriate assertability is preserved by arguments which would be valid on any acceptable precisifications of their vague terms, then it is clear why those natural language arguments which we treat as sound (on due reflection and adequate information) are useful to us, despite their not strictly speaking being sound. For we are led from sentences which engender in us appropriate expectations, ones which will not be frustrated, to others which will engender in us likewise appropriate expectations. Thus, on the assumption that appropriate assertability is preserved by those natural language arguments we think of as good arguments, we have an explanation of their utility despite their not being sound (nor unsound). They preserve in fact what is useful in asserting sentences containing vague predicates.

But why should we think that appropriate assertability in this sense is preserved in the cases in which we say it is? The reason is that those arguments

we treat as sound (on due reflection, etc.) are arguments that will be valid on any precisification of their terms that does not fly in the face of our practice. Suppose for a *reductio* that such an argument could have premises that engender appropriate expectations but a conclusion that does not. Since any acceptable precisification of its terms would have to be compatible with the practice, if the practice engenders expectations with respect to the assertion of the conclusion that are frustrated, any precisification would have to be false. But some precisifications of the premises will be true since those do not frustrate expectations. But the argument we have said is valid on all precisifications, contrary to this conclusion. Thus, the supposition we introduced must be rejected, and we can conclude that if the premises are appropriately assertable, so will the conclusion be, provided the argument is valid on all acceptable precisifications.

It will probably not have escaped the reader's notice that our characterization of 'appropriate assertability' used vague terms. In doing so, we have certainly engaged in the very practice that we aim to describe, that is, we have used vague terms to convey some information about the effect of using vague terms in sentences. There is nothing viciously circular about this, however. And the reader we have no doubt has understood us perfectly well. This is one place where using vague terms has particular utility. For it seems clear that there is in fact nothing both precise and general to say about what expectations are engendered in speakers by hearing sentences containing vague terms, and striving for something both precise and general is certain to yield something false. We know that in hearing a sentence containing a vague predicate asserted, such as 'Bert is bald', an auditor will in fact come to have some expectations about Bert. There will be a definite fact of the matter about whether the auditor will be surprised on seeing Bert or not, for example, in the light of the auditor's trust in the speaker and the circumstances of assertion. But can we really assign subjective probabilities for the auditor for each of the possible members of a sorites case involving Bert? It seems rather unlikely. Rather, we ourselves have certain general expectations about the auditor's expectations, but there is nothing general we can say about what to expect those to be for any individual, and it doesn't matter. What is key to understanding the effectiveness of the use of vague predicates is realizing that some expectations are set up which are conditioned by knowledge of the practice and trust in the speaker, and that if these expectations are close enough to those the speaker intends to induce, then the communicative exchange will be successful.³⁶

Some will no doubt be disturbed by reflection on our habit of using vague terms in sentences in the complement clauses of attitude reports. Some might be encouraged to think that because we use sentences with vague predicates in the complement clauses of attitude sentences, we are committed to there being vague thought contents and vague concepts (see (Machina 1976), for example). The cost of thinking this is high. If we are right about the non-truth-evaluability of sentences containing vague predicates, the same lesson will carry over straightforwardly to attitude contents, if we can conceive of such a thing, which are correspondingly vague: they will be neither true nor false. It is ques-

tionable whether we can make sense of belief that is not either true or false. Corresponding remarks apply to other attitudes. Fortunately, we are not forced to this extremity. For the tempting inference is an instance of what Russell called the fallacy of verbalism, which “consists in mistaking the properties of words for the properties of things” (Russell 1923, p. 147). Russell had in mind the inference from vague terms to vague objects, but the point applies equally to the inference from vague complement sentences to vague thought contents. On our account, what we do in such a case is no more mysterious than in any use of a vague predicate. Instead of applying a vague predicate to an object, we employ it in a sentence used putatively to give the content of an attitude. We fail to say something true, but we convey useful information nonetheless.³⁷ In general, we need no more fine-grained a picture of someone’s attitudes than we need of our environments. The fact that there are vague terms in our languages which we perforce employ in sentences characterizing attitude contents gives us no more reason to think our concept of a thought is of something with a vague content or that could have a vague content (something we frankly do not find fully intelligible) than it does to think that our logic must not be bivalent because not all of our sentences are fully meaningful.

It would not do, though, to leave the impression that there are no challenges here for us. We take it that thought contents are not vague, indeed, that it does not make sense to think of them so. And our conception of how vague terms are introduced is articulated against the backdrop of a conception of thought as not likewise vague, for we wish to say that we can explain how the phenomenon of vagueness comes about by appeal to intentions with regard to the use of terms not fully determining extensions for those terms. The challenge is to show that we can tell a story that makes sense of this in light of the difficulties presented by the pervasiveness of vague terms in natural language and our forced use of them in lieu of any precise replacement to describe our attitude contents. This is evidently an important challenge. We do not, however, undertake to meet it in the present context. What we think we need is a deeper understanding of the primitive sources of representation that underlie the introduction of linguistic representations. Achieving this understanding and applying it to the present discussion is too large a task to undertake here.

Let us return finally to our own free use of vague terms in this paper. While we do not endorse “the words of the poet: ‘Who speaks of vagueness should himself be vague’,”³⁸ it is certainly difficult to avoid the use of vague terms in discussing any subject in a natural language. As Michael Dummett has aptly said, “The great difficulty in discussing vagueness is that vagueness resembles dust, soot or sand: It gets into everything” (Dummett, 1995, p. 207). Our view is that there is a kind of talk involving vague words which is appropriate to our practice, but which we must carefully eschew if we want to get clear about the vagueness phenomena. So, with respect to a standard class of predicates that have been our focus of attention, we have refrained from engaging in our usual practice with these words. This has forced us to say things in certain ways that, ultimately, are conducive to getting clear about vague predicates. However, this

is not to say that we have not engaged in our ordinary communicative practices with respect to other terms, ones which were not our chosen subjects of discussion. So, with respect to a family of “front line” vague predicates we have refrained from the practice of the use of these terms, but in the process of our theorizing, we have employed certain “second-tier” vague terms. In using these terms we have just been speaking normally—participating in the usual language game.

Nonetheless, we have heard from an epistemicist³⁹ the objection that he can have understood nothing of what we’ve said, on the assumption that what we say is true. We are confident that this is not so, for we are right, and he did understand us, as his objection shows. But in light of the charge it is perhaps worthwhile pointing out that neither someone who takes the middle position nor an epistemicist is in any better position than we are in this respect. First, the epistemicist of course must deny that we know what propositions our sentences express. And the defect in our understanding on this view mirrors exactly what we think of as the defect in our semantics for vague predicates. Thus, the epistemicist is open to the charge equally that if he is right we do not know what he is saying. Of course, it is open to epistemicists to say that our ignorance is not complete, and that we know something about the likelihood of the sentences we use expressing one or another among the many propositions compatible with the vague grasp we have on what the sentences we use express. Thus, we can communicate information using sentences even if we do not understand what propositions they express. We would not wish to deny this, but only to insist that if this response is open to the epistemicist, a parallel response is open to us: our practices do put constraints of a sort on how the world can be, given an appropriately assertable sentence containing one or more vague predicates. This enables us to communicate information. We are in no better or worse position to make this point than is the epistemicist, so there is no comfort for him in raising an objection on this point. Second, someone who admits that vague terms are semantically defective but wishes to say some uses are in true sentences, some in false, and some in sentences that are neither true nor false has already admitted that none of these sentences are semantically complete, and, hence, none of them express propositions. The middle position theorist then is in exactly the same position that we are with respect to this objection.

IX

To summarize, we have argued that the default position on vague terms, namely, that they are semantically vague, is the correct position. It would be astonishing if this were not true. This fact poses no threat to the thought that the logic of natural languages is classical logic, i.e., that the semantics of natural languages and the characterization of its logical terms assigns to well formed semantically complete sentences either truth or falsity (ignoring complications of context sensitivity). That there are some vague terms in a language does not show that all are, and gives no support to the thought that we must accept a

non-classical logic. Nor does it pose any threat to the thought that we can use vague terms in our ordinary language practices to good effect. Vagueness is no mystery, or, at least, should not be. It does not force us to think either things or thoughts are vague because we happen to use vague terms to describe them. Our language concepts are simple; our language practices complex. The sorites arguments are, of course, not sound, but not unsound either. Their appeal and paradoxicality is explicable in terms of our default position. The minor premise is well within our practices in using the vague term in question; the conclusion well without. The step premise which causes the difficulty seems compelling not because it is well within our practice in using vague terms, but because it is a misbegotten material mode attempt at expressing a formal mode truth about the way in which our practices in the use of vague terms give out when we come to incremental changes along relevant dimensions of variation. We have argued that if one accepts that vague terms are semantically vague, then vague terms apply to nothing and fail-to-apply to nothing. There is no middle ground here. Either one accepts our result or denies boundarylessness by asserting there are first-order or second-order boundaries. If one rejects boundarylessness, then one is faced with the unenviable task of explaining why we cannot know the borderlines even in principle. We have argued that our position is ontologically neutral, and does not force us to deny anything we would otherwise assert in participating in our usual practices. We have argued vagueness engenders no difficulties for our ability to communicate information using vague terms, and does not undermine our ordinary argumentative practices under the pretense that our terms are not semantically vague.⁴⁰

APPENDIX

The appendix presents a proof of THEOREM given in section VII. THEOREM* follows from THEOREM, and has the conclusion we draw with the aid of THEOREM as its consequent.

THEOREM:

- IF (i) for all n , 'bald' applies to n iff 'bald' applies to $n + 1$, and
(ii) for all n , 'bald' fails-to-apply to n iff 'bald' fails-to-apply to $n + 1$,
THEN (iii) 'bald' applies to every n ,
(iv) 'bald' fails-to-apply to every n , or
(v) 'bald' neither applies nor fails-to-apply to any n .

Proof:

- AD For all n , if 'bald' applies to $n + 1$ then 'bald' applies to n . [given]
FU For all n , if 'bald' fails-to-apply to n then 'bald' fails-to-apply to $n + 1$. [given]
AU For all n , if 'bald' applies to n then 'bald' applies to $n + 1$. [given]
FD For all n , if 'bald' fails-to-apply to $n + 1$ then 'bald' fails-to-apply to n . [given]
1 Suppose 'bald' applies to some m . [supp]

- 2 'bald' applies to every $n < m$. [from 1, AD]
 3 'bald' applies to every $n > m$. [from 1, AU]
 4 'bald' applies to every n . [from 1,2,3]
 5 If 'bald' applies to some n , then 'bald' applies to every n . [from 1–4]
 6 Suppose 'bald' fails-to-apply to some m . [supp]
 7 'bald' fails-to-apply to every $n > m$. [from 6, FU]
 8 'bald' fails-to-apply to every $n < m$. [from 6, FD]
 9 'bald' fails-to-apply to every n . [from 6,7,8]
 10 If 'bald' fails-to-apply to some n , then 'bald' fails-to-apply to every n . [from 6–9]
 V Either 'bald' applies to some n , 'bald' fails-to-apply to some n , or 'bald' neither applies nor fails-to-apply to any n . [truth of logic]
 C Either 'bald' applies to every n , 'bald' fails-to-apply to every n , or 'bald' neither applies nor fails-to-apply to any n . [from 5,10,V]

THEOREM*:

- IF (i) for all n , 'bald' applies to n iff 'bald' applies to $n + 1$, and
 (ii) for all n , 'bald' fails-to-apply to n iff 'bald' fails-to-apply to $n + 1$,
 (iii) there is some n such that 'bald' does not apply n ,
 (iv) there is some n such that 'bald' does not fail-to-apply n ,
 THEN (v) 'bald' neither applies nor fails-to-apply to any n .

Notes

1. See, for example, (Parsons 1987), (Tye 1990), (Zemach 1991), (Parsons and Woodruff 1995).
2. Epistemicism is said to originate with the Stoics. It was reintroduced this century by (Cargile 1969). The most sustained defense of this position is (Williamson 1994). Other defenders are (Campbell 1974) and (Sorenson 1988).
3. See Sainsbury (1991, 1999) for a discussion of this feature of vague predicates, to whom we owe this apt phrase. (See Horgan (1990, 1994) for similar arguments.) Boundarylessness entails that there are no higher-order borderlines as well as no first-order borderlines. There are no borderlines at all for boundaryless predicates. Often the phenomenon of vagueness is discussed under a certain idealization, that there are clear positive cases, clear negative cases, and borderline cases of vague predicates—and that all cases fall neatly into one of these categories. Supervaluational theories, for example, presuppose this. But there is no clear dividing line between cases we intuitively feel it is acceptable to use a vague predicate with respect to and those we feel uncertain about. Imagine (what we would call) a heap of finely ground white pepper a grain of which is removed at a time. It is clear enough we will not be able to say at which point the removal of a grain of pepper results in our moving from a positive case for the use of "is a heap" to a borderline case. The same applies to any higher-order borderline. This is the central feature of the "semantic profile" of the predicates we are interested in.
4. We offer a general argument for the general claim and a proof in the case of vague predicates that exhibit the feature of boundarylessness.

5. The modern *locus classicus* of discussion of vagueness is (Russell 1923). While Russell was on the right track, he sets the tone for most discussion this century by assuming some applications of vague predicates express truths, some falsehoods, and some neither. All varieties of many valued logics, including supervaluationist views, when treated as capturing semantic facts about uses of natural language sentences, likewise accept this assumption. (Hallden 1949) is an early example of the application of a three-valued logic to the semantics of vague sentences. (Fine 1975) is the classical discussion of supervaluationism. Representatives of degrees of truth and fuzzy logic approaches are (Goguen 1969) and (Machina 1972). See also more recently (Forbes 1983; Horgan 1994; Horgan 1995).
6. *Pace* Dummett (1995, p. 211), this is not to abandon a two-valued semantics. For being neither true nor false is not *another semantic value* a sentence may have analogous to truth and falsity. It is the sentence's failing, like a horse, to be something that qualifies for having a semantic value.
7. We have heard it objected that this need not be a defect, since having semantically incomplete terms in natural language may be just what we need given the exigencies of our practices. This is like saying that having a beat up car with missing hubcaps and a rusting trunk lid need not be considered a defect if one has to park it on the streets in Manhattan. These are defects, but useful in some circumstances. Likewise, semantic defects can sometimes be useful. Even nonsense words have their utility. (They make good passwords, for example.) But it doesn't follow that semantic incompleteness is not a defect relative to the presuppositions of our semantic predicates.
8. Semantic completeness is not the same thing as its following from the facts about a practice whether or not a term applies, fails-to-apply, or neither applies or nor fails-to-apply. A term is semantically complete if the aggregate intentions (and other relevant facts) with respect to its use determine for any case whether it applies, fails-to-apply, or neither. But if the meaning determining practice is radically defective, as it is with nonsense terms that are accorded merely a grammatical category, it can hardly be said that the term is semantically complete, though the facts about the practice, namely, that it does not determine any applicability conditions, determine that the nonsense term does not apply or fail-to-apply to anything. Unger's argument against the possibility of incomplete terms in (1979b) comes to grief on failure to note this distinction. It may be that Unger's failure likewise to distinguish *not applying* from *failing-to-apply* helped him along, since the conclusion might seem to follow if one failed to distinguish not applying from failing-to-apply and, from concluding that it followed from the facts about the practice that a term did not apply, concluded that it failed-to-apply, which would require the practice's determining in our sense semantic completeness.
9. There presumably will be (waiving concerns about physical indeterminacy) a definite fact about such things as what the number of grains is of a certain size in a certain arrangement that is the average of the number at which actual English speakers (would) withhold application of 'heap' (fixing enough conditions). This does not mean that our practices determine a rule that says that fewer than that number (in that configuration) no longer determines a heap. For *there would be such a fact whether or not our meaning-determining practices left it open*.
10. When we speak of practices in contexts in which it is clear we are concerned with what rules govern the applicability conditions of a term, we mean meaning-determining practices; henceforth we drop the modifier.

11. Compare (Quine, 1985, p.168):

Are we then to withhold the term ‘physical object’ from the very things that have been its prototypes—desks and mountains? Yes and no. A certain adjustment is required, and the place where I would make it is in the interval between formal logic and the terms to which it is applied. Consider, to begin with, the classical notion of the extension of a general term. The extension of the term ‘desk’ is conventionally thought of as the class of its denotata, thought of as physical objects. Realistically we may recognize rather an extension family, as I shall call it. It is a family of vaguely delimited classes, each class being comprised of nested physical objects any of which would pass indifferently for one and the same desk. When we bring formal logic to bear on discourse of desks, then, we adopt the fiction that the extension is some one arbitrary and unspecified selection class from that family of classes; it selects one physical object from each. Similarly, and more obviously perhaps, for mountains. This strikes me as the reasonable way to accommodate vagueness: not in a logic of vagueness, but in the account of the application of a logic of precision.

12. It may well be worth noting at this point that this does not commit us to the material mode formulation of the position which Timothy Williamson (1994) has called ‘nihilism’, that, e.g., nothing is a heap, no one is bald, etc. Vague predicates neither apply nor fail-to-apply to anything. We could infer ‘No one is bald’ from ‘“bald” fails-to-apply to anyone’, but not from ‘“bald” neither applies nor fails-to-apply to anyone’. This is one point at which we can see how confusion can arise from failing to attend to our three-fold distinction. We will return to this point when we discuss Unger’s arguments from vagueness to negative ontological conclusions in section VIII.
13. Our discussion focuses on indicatives, but the remarks extend with appropriate adjustments to imperatives and interrogatives.
14. Crispin Wright (1976, p. 223) says, “... Frege seems not to have considered, or not to have thought worth considering, the possibility that vague terms might require a *special* logic.” What is meant here by ‘a special logic’? It is not explained. Is it a special set of non-classical logical terms? But then why should Frege think vague terms required their own special set of logical terms? And required in what sense?
15. It is surprising how routinely this distinction is overlooked. Supervaluationism and virtually every multi-valued approach to vagueness fail to mark it. A particularly clear early example is provided by a paper by Max Black (1937), in which what is essentially a model of speakers’ dispositions to apply or withhold a vague term is taken to provide a guide to the semantics of the language. One would think that vagueness had a strange power to cloud the mind of anyone who turned his attention to it. Hempel (1939), responding to Black, is a notable exception.
16. We will suppose this is so despite Lewis Carroll’s having something in mind. The stanza of “Jabberwocky” from which this is taken appeared originally in *Mischmasch*, a magazine written by and for the Carroll family in 1855, when Carroll was 23. It originally read “And ye mome raths outgrabe,” and Carroll explained that it meant “and the grave turtles squeaked out” (Carroll 1960, p. 192). He apparently changed his mind when he included this stanza in “Jabberwocky” in *Through the Looking Glass*. The relevant words there receive the following explanation:

“And then ‘mome raths’?” said Alice. “I’m afraid I’m giving you a great deal of trouble.”

“Well, a ‘rath’ is a sort of green pig; but ‘mome’ I’m not certain about. I think it’s short for ‘from home’—meaning that they’d lost their way, you know.”

“And what does ‘outgrabe’ mean?”

“Well, ‘outgrabing’ is something between bellowing and whistling, with a kind of sneeze in the middle: however, you’ll hear it done, maybe—down in the wood yonder—and, when you’ve once heard it, you’ll be quite content.”
(Carroll 1960, pp. 270–2)

Despite this, we’re inclined to say that *in English* these are nonsense words. In any case, the point is unaffected.

17. Suppose we introduce a function defined over the real numbers that takes a pair of real numbers to their sum: $f(x, y) = x + y$. Suppose someone asks, what is the value of the function for $\langle \text{Bush}, \text{Cheney} \rangle$? The answer is nothing. Let us then represent this as ‘N’, and give the following table to represent the facts where ‘ r_1 ’ and ‘ r_2 ’ represent real numbers.

x	y	$f(x, y)$
r_1	r_2	$r_1 + r_2$
Bush	Cheney	N

Now suppose someone staring at this table mistakes it for, or confuses it with, a definition of the function with ‘N’ representing yet another value the function has for the pair $\langle \text{Bush}, \text{Cheney} \rangle$. We would have here a mistake analogous to that made in thinking that sentences are three-valued because they aren’t true or false if they contain nonsense words (*mutatis mutandis* for various “multi-valued” models of vagueness). The mistake would be exactly the same if we adopted a Fregean view according to which the sentential connectives expressed functions from the truth-values, the True and the False, to one or the other.

18. For our purposes, it is enough to identify expressing a complete sense with being grammatically well-formed from semantically complete terms for a non-context-sensitive declarative sentence, and being grammatically well-formed from semantically complete terms interpreted relative to contextual parameters that fix the contributions of its context sensitive elements for a context-sensitive declarative sentence. Grammatical well-formedness here will require attention to the appropriate categories of terms, so that, e.g., ‘The Milky Way galaxy is divisible by 2’ will not count as well-formed. Note that a predicate whose rules give rise to a contradiction will fail to have an extension, and so will not count as semantically complete, though, so to speak, through being over endowed rather than under endowed with meaning.
19. Discussion of the principle of bivalence is supposed to trace back at least to Aristotle and his discussion of future contingents. (BV) is a reasonable (sentential) version of the principle that exercised Aristotle in those passages. Future contingent sentences express complete senses, but it was at issue whether they had truth-

values since which value they would have to have evidently depended on things yet to happen.

20. Our use of vague predicates is tolerant with respect to small variations along relevant dimensions. The key point we urge is that this should not be misunderstood as a license on the part of our practices to apply the term to any case that differs only incrementally along a relevant dimension of variation if it applies to the given case. This would make the practice incoherent, rather than simply incomplete. It is failure to distinguish the practice giving no guidance from its providing a license to move from one case to an incrementally different one which is the source of much of the confusion about the sorites arguments. Sainsbury (1999, p. 260) makes a related point, distinguishing between its not being mandatory to withhold a vague predicate from one of two members of pair that differ incrementally if the predicate is applied to the other and its being mandatory not to withhold it. The latter would express a positive rule; the former can be explained by the practice's simply giving no guidance. Peter Unger's arguments in (1979a) and (1979b), for example, come to grief on failure to mark this distinction. Unger argues from tolerance to the step premise, then runs the sorites arguments in both directions, and concludes that since a contradiction results, the term is "inconsistent," and nothing falls in the extension of any such predicate. Michael Dummett (1975) also takes tolerance to sanction the step premise, and concludes likewise that this means our practice is incoherent. Dummett singles out observational predicates for special attention, arguing that it is a rule for their use that if they apply to a thing they apply to a thing indiscriminable from it, so that this together with the non-transitivity of perceptual discrimination leads to paradox. But he extends the claim to all vague predicates, whether observational or not. See note 34 for further discussion. If we are right, these reactions are unwarranted. Each fails to take seriously enough the thought that the practice attached to these predicates simply does not answer every question we might have about their applicability conditions. (Crispin Wright has used the term 'tolerance' to describe predicates as having attached to them a positive rule to the effect that application to one item is licensed if it differs relevantly incrementally from another to which application is licensed (1976). Wright does not endorse this view, rather, he argues against what he calls the governing view of language in part on the grounds that it is committed to taking the step premise to express a positive rule, and so landing us with predicates governed by rules which sanction the derivation of contradictions. Again, see note 34 for relevant discussion.)
21. The shift to the formal mode is essential in getting clear about the appeal and fallaciousness of the sorites reasoning. Consider Galen's recounting of the inevitability of the step premise in the argument of the heap: "And I know of nothing worse and more absurd than that the being and not-being of a heap is determined by a grain of corn" (*On Medical Experience* 16.1–17.3 translated in (Long and Sedley 1987, 222–25)). To use a term to express something is to engage in a practice that presupposes the term is semantically complete. If one overlooks the distinction between using a term and talking about it, in trying to express the fact that our practice does not say that individual grains make the difference between whether 'heap' is to apply or fail-to-apply, one will inevitably express it ontically, and then all is lost, for this is in effect to take the silence of our practice for positive advice. As Russell says, "The influence of symbolism on philosophy is mainly unconscious; if it were conscious it would do less harm" (1923, p. 84).

22. A sorites argument can also be formulated using the minor premise and a series of conditionals which are of the form ‘if a man with N hairs on his head is bald, then a man with $N + 1$ hairs on his head is bald’. The diagnosis goes the same way here. These material mode conditionals we are drawn to accept because they seem, as a group, expressive of the vagueness of ‘bald’. But this is to mistake no instruction about incremental changes in the number of hairs on someone’s head for a positive rule. Our usual practice works because we avoid using the predicate in circumstances in which the practice gives us no guide. The sorites arguments generate trouble because they invite us to engage in our usual pretense about completeness, but in areas where the pretense breaks down.
23. If ‘bald’ neither applies nor fails-to-apply to any x , then ‘applies’ and ‘fails-to-apply’ fail-to-apply to \langle ‘bald’, x \rangle for any x . ‘applies’ then applies to \langle ‘fails-to-apply’, \langle ‘applies’, \langle ‘bald’, x \rangle \rangle , and to \langle ‘fails-to-apply’, \langle ‘fails-to-apply’, \langle ‘bald’, x \rangle \rangle , for any x , and so on.
24. Russell makes this mistake in his (1923). Russell argued that logical connectives in natural languages must be vague because they were defined in terms of the semantic predicates ‘true’ and ‘false’ ($\langle p \text{ or } q \rangle$ is true if ‘ p ’ is true or ‘ q ’ is true and otherwise false), and the semantic predicates were perforce vague in application to a vague language. Russell says, “Now ‘true’ and ‘false’ can only have a *precise* meaning when the symbols employed—words, perceptions, images, or what not—are themselves precise. We have seen that, in practice, this is not the case” (p. 64). And

... it is possible to discover prehistoric specimens concerning which there is not, even in theory, a definite answer to the question, “Is this a man?” As applied to such specimens, the proposition “this is a man” is neither definitely true nor definitely false. Since all non-logical words have this kind of vagueness, it follows that the conceptions of truth and falsehood, as applied to propositions composed of or containing non-logical words, are themselves more or less vague. (p. 65)

The argument intended here seems to be the following (we suppress relativization to a language).

- (i) There is a range of objects m_1, \dots, m_j such that it becomes more and more difficult to answer the question ‘Is m_i a man?’ as one moves from m_1 to m_j , (though the fault lies not in any failure in our knowledge of how the relevant portions of the world are or the meanings of the terms in the sentence—assume this qualification repeated below).
- (ii) To the degree that it is difficult to answer the question ‘Is m_j a man?’, it is to the same degree difficult to say whether ‘ m_j is a man’ is true (or false).
- (iii) Therefore, there is a series of sentences, ‘ m_1 is a man’, ‘ m_2 is a man’, ... ‘ m_j is a man’, to which the application of ‘is true’ (‘is false’) is doubtful in the same way and to the same degree as is the application of ‘is a man’ to $m_1 \dots m_j$. [(i), (ii)]
- (iv) Therefore, ‘is true’ (‘is false’) is vague in the same way and to the same degree as ‘is a man’. [(iii)]

If what we have said in the text is correct, the appearance that (ii) is correct is an illusion. It is difficult only so long as we continue to labor under the usual pretense that ‘is a man’ is a precise predicate: give this up, and the answer is clearly and determinately that no sentence in the series ‘ m_1 is a man’, ... is true or false.

25. So far as we can see, no precisification of vague semantic predicates could fail to place all pairs of vague terms and objects in the third category. But suppose that this is not true (somehow). Still it will be admitted that the argument for the THEOREM will be sound on any precisification of the terms in it. If an argument sound on all of its precisifications is to be thought a good one, then even someone who wants to insist natural language semantic terms are vague should admit that the argument is a good one.
26. If anyone doubts there are or can be such series, we refer him to the opening paragraph of Cargile's "The Sorites Paradox" (1969). Moreover, such series, of course, need not be of spatio-temporal objects. A sorites argument can be run on the predicate 'is a large number'.
27. We have carefully not characterized boundarylessness in terms of iterations of the operator 'it is indefinite that'. How is such an operator to be explained? What is indefinite and in what sense? It is difficult to see what could be intended other than 'it is indefinite whether it is true or false that', where it is taken not to mean that it is definitely neither, but rather it is in some sense undetermined. (It might be maintained to be primitive—in that case *we* do not understand sentences containing it, which is a good enough reason not to use it in characterizing anything.) And what is it that is at issue? A sentence or a proposition? On our view, according to which it is never undetermined whether a sentence is true, false, or neither true nor false, if sentences are the object, no application of the operator yields a true sentence. If propositions are said to be the object, at best an application could yield a false sentence when applied to a sentence that contains no vague terms. When applied to a sentence containing vague terms, the *resulting* sentence would be in our view neither true nor false, just as the contained sentence is. There is then no room for this operator in our account. If we are right, presuppositions are introduced with these operators that have made it difficult to see the matter clearly. If we interpret 'it is indefinite that' simply as 'it is neither true nor false that', then if we take expressions that come after this operator to be used in the sentence in a way that requires predicates to contribute their extensions to fixing truth conditions, then again no sentence of the form 'it is indefinite that ϕ ' will be true or false if ϕ contains a vague predicate, and neither will any sentence in which this is used, including sentences formed from it by iterating the operator. If we take the position of ϕ to be metalinguistic, on the other hand, then every sentence of the form 'it is indefinite that ϕ ' will be either true or false, and so likewise any sentences formed from it by one or more applications of the same operator.
28. A note is perhaps in order about the so-called contextualist response to sorites arguments. (See (Raffman 1994) for an example.) These accounts are most popular when the problem of vagueness is confounded with the failure of transitivity of perceptual discrimination, and so are typically advanced when an author is talking about the vagueness of color terms or other terms picking out determinates of perceptual determinables. Contextualists suggest that what appears to be semantic vagueness can be diagnosed as semantic context sensitivity of the relevant terms and our failure to attend to shifts in context. No one has advanced any clear account of what rules would attach to our predicates that would determine relative to any context what the borderline for that context would be. This would have to take the following form:

For any contextual parameters c_1, \dots, c_n , any speaker S and time t , and any object x , ζ applies in L to x relative to S at t and c_1, \dots, c_n iff $Z(x, S, t, c_1, \dots, c_n)$,

where ‘ $Z(x, S, t, c_1, \dots, c_n)$ ’ stands in for a precise metalanguage predicate with argument places for object, speaker, time, and each relevant contextual parameter. Does anyone pretend to be in a position for, say, ‘bald’, to say what is to go in for ‘ $Z(x, S, t, c_1, \dots, c_n)$ ’ that determines relative to any context exactly the breakpoint number of hairs? Of course not, and this is unsurprising since it is clear that we don’t understand there to be such a rule. Take a particular concrete context in which we are faced with a man with, oh, just so much hair on his head—one of the puzzling cases. Now that it is clear that we have fixed the context, is it clear exactly how many hairs it takes for someone relative to that context to be bald, and so, after we ascertain the facts, whether this man is bald? Not at all! It is no help. Even if these terms exhibit some kind of context relative interpretation, which we think is doubtful, it is clear enough it does not eliminate the phenomenon. The reason these suggestions arise in the case of phenomenal sorites arguments is that there is some work for context to do here, and that is in making sense of the possibility of failure of transitivity of indiscriminability. How can A appear indiscriminable from B with respect to F and B from C for S if A and C are discriminable? Would S not then have inconsistent perceptual representations of the world? How are we to make sense of that, in, e.g., the case of color? The answer is to relativize the perceptual judgments of indiscriminability to temporally indexed contexts. There is no difficulty about saying that at t A appears indiscriminable with respect to F from B for S, at $t' \neq t$ B appears indiscriminable from C, and at $t'' \neq t'$ or t A appears discriminable from C for S. At each time our perceptual representations are consistent. Across time they are not. But the latter entails no single phenomenal experience whose content is inconsistent.

29. One move here might be to urge that all our vague terms are like natural kind terms, and have their extensions fixed in some complicated way by the things to which we actually apply them. We have a model of this in the case of a term like ‘gold’ or ‘water’. The extension is that set of things which are alike with respect to the property F of kind K which is such that most members of the set of extension-fixing paradigms have F and their having F explains their having the phenomenal properties by which we identify them. The kind K is determined by the kind of theory appropriate for the kind term in question. For biological kinds, it would be biological theory—genetic code; for kinds like ‘gold’ or ‘water’, the relevant theory is chemistry, and we look for some chemical structure. This provides us with a means of making sense of how we could discover the extensions by empirical investigation. Could this model be used for vague terms generally? It is not easy to see how. Consider ‘table’. What we would need is some way of making sense of a relevant kind K which determines a range of properties from which the way the world actually is selects an appropriate one which determines where the borderlines are for all the dimensions of variation relevant in the practice for the use of the term. What could this be? How plausible is it that we intend such terms to be hostage to how certain empirical investigations turn out? Where we have no idea what empirical investigations would be relevant, this is a sign that it is not part of our practice to hold the extensions of such terms hostage to empirical investigations.
30. The most elaborate attempt is surely Timothy Williamson’s in *Vagueness* (1994). However, at best his solution, not itself unproblematic, would yield contingent ignorance of borderlines, not necessary ignorance. See Ray (1999) for a critique of Williamson’s argument.
31. See section 5 of Horgan’s (1994) for further criticism of the hidden facts theory.

32. See (Williamson 1992) and (Williamson 1994) for a forceful expression of this view.
33. Our remarks apply *mutatis mutandis* to non-declaratives of course. For a generalization of the truth-theoretic approach to semantics to non-declaratives, see (Ludwig 1997).
34. As has been observed by (Dummett 1975), the non-transitivity of perceptual indiscriminability for many observational properties also makes vague terms a natural fit for us. If we cannot tell differences when they are small enough without special tests, then there will be no point in trying to make our practices sensitive to such differences for ordinary purposes. Dummett suggests that the non-transitivity of indiscriminability will inevitably lead us into troubles because we operate using observational predicates with the assumption that if a predicate applies to a thing it applies to a thing indiscriminably different. We don't see, however, why this should be thought to be a rule of meaning determined by our practices. We recognize in general the possibility of a gap between perceptual representation and reality. This is enough to allow that there may be an indiscriminable difference between things. It would be strange if our meaning determining practices were at odds with this. In cases in which applying such a rule would get us into trouble we can in fact deduce that there are indiscriminable differences. If A and B are indiscriminable with respect to color and B and C are also, but A and C are discriminable, we can infer A and B and B and C differ, though not discriminably so. That our concepts even of observables allow this degree of objectivity to the properties we take ourselves to perceive is supported by our recognition that individuals can differ in discriminatory abilities. Sue may be able to distinguish tones or tastes Abe cannot. We do not think Abe and Sue inhabit different worlds. Sue has access to properties Abe does not. So likewise we can make room for the thought that there may be differences in determinates of observable properties that none of us can discriminate. This is of a piece with thinking of the properties we have access to in perception as objective. Notice furthermore that this idea presents no difficulties to us when we are considering comparatives such as 'x is more violet than y'. Given the evidence about discriminability of A, B and C above, and assuming C looks more violet than A, we would straightforwardly conclude that B is more violet than A and less violet than C. There is no conflict with our concepts of color properties in this judgment. All of this is, furthermore, compatible with there being purely observational predicates, *pace* (Wright 1976). A predicate is purely observational if one can tell it applies to a thing by looking. Failure of transitivity of sameness of appearance looks like a threat to this, since it is natural to think that it requires us to say any two things that look alike in relevant respects when compared receive the same designation. But it does not require this. A predicate can be purely observational if the requirement on its application is that it look the same with respect to F (some determinable property) as some paradigm, and one can acquire a disposition to apply the predicate on the basis of whatever mechanism is in place which yields judgments of sameness of F when explicit comparisons are made. We can see no reason why this should be impossible. In this case, the application of the predicate need not involve any explicit comparison with the paradigm, of course, though our exposure to it guides us in evaluating its look. Something *x* may then fall under it which when compared with another thing *y* looks to be the same F, though *y* does not fall under it, because our judgment about *y* activated by sensitivity to the determinate property of the paradigm is that *y* is not the same with respect to F. The point is that whatever mechanism underlies the difference in judgments in cases of explicit comparison

can be internalized to provide competence in the application of a predicate directly on observation.

35. An acceptable precisification of a set of terms must among other things respect what Kit Fine has called ‘penumbral connections’, such as, e.g., expressed by our practice of withholding ‘red’ from anything we are willing to call ‘orange’ and vice versa. This is of a piece with our treating these different predicates as expressing distinct determinates of a determinable. The notion of an acceptable precisification is itself vague, though it may be that it could be given a purely behavioral sense that was precise. We say something to ameliorate worries about this below.
36. In our view, most formal treatments of vagueness masquerading as semantics for vagueness are better understood as aiming at modeling something like appropriate use, or something in the area. Typically, however, the models read too much structure into the use. The trouble is that our practices under specify use. They don’t specify degrees of appropriateness of use of any sort that could be captured correctly even by a continuum of values. Edgington (1999) notes this explicitly and makes a good case for the utility of such models even so, though she is not entirely on board with our way of thinking about vagueness.

Edgington develops an account in terms of degrees of truth, or verities, to use her helpful technical term, which models the semantic facts about vague sentences on analogy with degrees of belief or credences in propositions. She aims to make sense in terms of this model of arguments being valid though their conclusions are not acceptable because their unverity is high while the unverity of their individual premises may not be anywhere near so high. We think much of what Edgington has to say can be usefully exploited with a slight reorientation for our purposes. On her account, a natural language argument can be seen to be valid iff the unverity (1—the verity of a proposition) of the conclusion cannot be higher than the sum of the unverities of the premises. The relation of the verities of molecular sentences to their components is modeled on the relations between the probabilities of molecular sentences and those of their components (this gives up the usual “truth”-functionality of degrees of truth models of vagueness—Edgington makes a convincing case for this). If we read ‘verity’ (under an idealization, of course) as degree of appropriate assertability, we can get a model that allows us to call all arguments valid while admitting that the conclusion may not be appropriately assertable. For if Edgington is correct, on this model, a sorites argument consisting of a minor premise and a series of conditionals yielding an unacceptable conclusion will be valid but this will be consistent with the conclusion receiving a low verity. (We must assume here a finite sorites series between end points whose verities are between 1 and 0.) For each conditional (taken to be equivalent to the negation of the conjunction of the antecedent and negation of the consequent) will have small unverity equal to the difference in degree of assertability of a sentence predicating the relevant predicate of an item of the first sort and the degree for a sentence predicating it of an item of the second sort, say, ϵ . Assuming the unverity of the first premise is zero, and the number of steps in the argument is n , the unverity of the conclusion must be less than or equal to $n \times \epsilon$. For large enough n , of course, no matter how small ϵ (assuming it divides the space between the end-points evenly), this can be close to 1, and correspondingly the verity can be quite low. Indeed, if we start at one endpoint of a series and move to the other, which is treated as having a verity of 0, the sum of the unverities will be 1. Edgington extends the point to the case involving a universally quantified step premise over

- the domain of the sorites series, more or less by fiat, arguing the verity of the quantified step premise must be zero for consistency.
37. If our own account of the semantics of attitude sentences is correct (Ludwig and Ray 1998), sentences attributing attitude reports will all be false, rather than neither true nor false (waiving worries about terms not in the complement clause). This is because we endorse a sophisticated sententialist account that requires sameness of content between complement clauses or completions of them interpreted relative to contexts and attitude states for the truth of an attitude report. Since the words in the complement clause, though they must be understood to understand the report, do not contribute to the truth conditions of the attitude sentences, their vagueness does not infect the truth conditions of the whole sentence.
 38. This is from Russell's (1923, p. 147), a literary joke on Russell's part, after Samuel Johnson's rejoinder, "you might as well say, 'Who drives fat oxen should himself be fat,'" to Henry Brooke's line in *Gustavus Vasa* (thought to be directed at Robert Walpole) "Who rules o'er free men should himself be free." We are indebted to Roland Hall for identifying the allusion. The annotation in (Slater 1988, p. 522) on this line cites a line from Henry Brooke's play *The Earl of Essex*, which is a variant of the line from *Gustavus Vasa*, "To rule o'er Freemen, should themselves be free."
 39. Eugene Mills raised this objection to us when an earlier version of this paper was presented at the Inter-University Center conference on Vagueness in Bled, Slovenia, June 1998.
 40. We would like to thank Patricia Blanchette, Eugene Mills, Piers Rawling, Cara Spencer, and audiences at the 1998 Bled Conference on Vagueness, the 1999 Pacific Division APA meetings, the 1999 Society for Exact Philosophy meetings, the 1999 Florida Philosophical Association meetings, the University of Missouri, St. Louis, Wayne State University, Notre Dame University, and Webster University for comments and suggestions on earlier versions of this paper.

References

- Black, Max. 1937. Vagueness: an exercise in logical analysis. *Philosophy of Science* 4:427–55.
- Campbell, R. 1974. The Sorites Paradox. *Philosophical Studies* 26:175–91.
- Cargile, James. 1969. The Sorites Paradox. *British Journal for the Philosophy of Science* 20:193–202.
- Carroll, Lewis. 1960. *The annotated Alice: Alice's adventures in Wonderland & Through the looking glass*. With an introduction and notes by Martin Gardner. New York: C. N. Potter.
- Dummett, Michael. 1975. Wang's Paradox. *Synthese* 30:301–24.
- Dummett, Michael. 1995. Bivalence and Vagueness. *Theoria* 61:201–216.
- Edgington, Dorothy. 1999. Vagueness by Degrees. In *Vagueness: A Reader*, edited by R. Keefe and P. Smith. Boston: MIT Press.
- Fine, Kit. 1975. Vagueness, Truth and Logic. *Synthese* 30:265–300.
- Forbes, Graham. 1983. Thisness and Vagueness. *Synthese* 54:235–59.
- Goguen, J. A. 1969. The Logic of Inexact Concepts. *Synthese* 19:325–73.
- Haack, Susan. 1980. Is Truth Flat or Bumpy? In *Prospects for Pragmatism*, edited by D. H. Mellor. Cambridge: Cambridge University Press.
- Halden, Soren. 1949. *The Logic of Nonsense*. Uppsala: Lunde Bok.
- Hempel, Carl. 1939. Vagueness and Logic. *Philosophy of Science* 6:163–80.
- Horgan, Terence. 1990. Psychologistic Semantics, Robust Vagueness, and the Philosophy of Language. In *Meanings and Prototypes: Studies in Linguistic Categorization*, edited by S. L. Toshohatzidis. London: Routledge.
- Horgan, Terence. 1994. Robust Vagueness and the Forced-March Sorites Paradox. In *Logic & Language*, edited by J. Tomberlin. Atascadero: Ridgeview.

- Horgan, Terence. 1995. Transvaluationism: A Dionysian Approach to Vagueness. *S J Phil* 33:97–126.
- Long, A. A., and D. N. Sedley. 1987. *The Hellenistic Philosophers*. Cambridge: Cambridge University Press.
- Ludwig, Kirk. 1996. Singular Thought and the Cartesian Theory of Mind. *Nous* 30 (4):434–60.
- Ludwig, Kirk. 1997. The Truth about Moods. *Protosociology* 10 Cognitive Semantics I—Conceptions of Meaning:19–66.
- Ludwig, Kirk, and Greg Ray. 1998. Semantics for Opaque Contexts. *Philosophical Perspectives* 12:141–166.
- Machina, Kenton F. 1972. Vague Predicates. *American Philosophical Quarterly* 9:225–233.
- Machina, Kenton F. 1976. Truth, Belief and Vagueness. *Journal of Philosophical Logic* 5:47–78.
- Naess, Arne. 1953. Interpretation and Preciseness: a contribution to the theory of communication. Oslo: I Kommissjon Hos J. Dybwad.
- Parsons, Terrence. 1987. Entities without Identity. In *Philosophical Perspectives, 1: Metaphysics*, edited by J. Tomberlin. Atascadero: Ridgeview.
- Parsons, Terrence, and Peter Woodruff. 1995. Worldly Indeterminacy of Identity. *Proceedings of the Aristotelian Society* 95:171–91.
- Quine, W. V. 1985. Reification and Events. In *Actions and Events: Perspectives on the Philosophy of Donald Davidson*, edited by E. Lepore & B. McLaughlin. Oxford: Basil Blackwell.
- Raffman, Diana. 1994. Vagueness without Paradox. *Philosophical Review* 103:41–74.
- Russell, Bertrand. 1923. Vagueness. *Australasian Journal of Philosophy* 1:84–92.
- Sainsbury, R. M. 1991. Is There Higher-Order Vagueness? *The Philosophical Quarterly* 41:167–182.
- Sainsbury, R. M. 1999. Concepts without Boundaries. In *Vagueness: A Reader*, edited by R. Keefe and P. Smith. Boston: MIT Press.
- Slater, John G. 1988. *Bertrand Russell: Essays on Language, Mind and Matter, 1919–26. The Collected Papers of Bertrand Russell*, 9. London: Unwin Hyman.
- Sorenson, Roy. 1988. *Blindspots*. Oxford: Clarendon Press.
- Tye, Michael. 1990. Vague Objects. *Mind* 99.
- Unger, Peter. 1979a. There Are No Ordinary Things. *Synthese* 41:117–54.
- Unger, Peter. 1979b. Why There Are No People. In *Studies in Metaphysics*, edited by P. A. French, T. E. Uehling and H. K. Wettstein. Minneapolis: University of Minnesota Press.
- Williamson, Timothy. 1992. Vagueness and Ignorance. *Proceedings of the Aristotelian Society, Supp.* 66:145–62.
- Williamson, Timothy. 1994. *Vagueness*. London: Routledge.
- Wright, Crispin. 1976. Language-mastery and the Sorites Paradox. In *Truth and Meaning*, edited by G. Evans and J. McDowell. Oxford: Oxford University Press.
- Zemach, Eddy. 1991. Vague Objects. *Nous* 25: 323–40.