PHILOSOPHICAL TOPICS

VAGUENESS

REFERENCES

Broome, John. 1984. "Indefiniteness in Identity." *Analysis* 44: 6–12.

Burgess, John. 1989. "Vague Identity: Evans Misrepresented." *Analysis* 49: 112–19.

———— 1990. "Vague Objects and Indeterminate Identity." *Philosophical Studies* 59: 263–87.

Edgington, Dorothy. 1996. "Vagueness by Degrees." In *Vagueness: A Reader*, ed. R. Keefe and P. Smith. Cambridge Mass.: MIT Press, 294–316.

Evans, Gareth. 1978. "Can There Be Vague Objects." *Analysis* 38: 208. Reprinted in his *Collected Papers*, 176–77. Oxford: Clarendon Press, 1985.

Fine, Kit. 1989. "The Problem of De Re Modality." In *Themes from Kaplan*, ed. J. Almog, J. Perry, and H. Wettstein. Oxford: Oxford University Press, 197–272.

Kaplan, David. 1986. "Opacity." In *The Philosophy of W. V. Quine*, 229–89, ed. L. E. Hahn and P. A. Schilpp. Chicago and La Salle, Ill.: Open Court.

Kripke, Saul. 1980. *Naming and Necessity*. Oxford: Basil Blackwell.

Lewis, David. 1988. "Vague Identity: Evans Misunderstood." *Analysis* 48: 128–30.

Lewis, David. 1993. "Many, but Almost One." In *Ontology, Causality and Mind: Essays on the Philosophy of D. M. Armstrong*, ed. K. Campbell, J. Bacon, and L. Reinhardt. Cambridge: Cambridge University Press. Reprinted in D. Lewis. *Papers in Metaphysics and Epistemology*, 164–82. Cambridge: Cambridge University Press, 1999.

Parsons, Terence, and Woodruff, Paul. 1995. "Worldly Indeterminacy of Identity." *Proceedings of the Aristotelian Society* 95: 171–91.

Priest, Graham. 1998. "Fuzzy Identity and Local Validity." *Monist* 81: 331–42.

Quine, W. V. O. 1960. *Word and Object*. Cambridge, Mass., MIT Press.

Richard, Mark. 1986: "Quantification and Leibniz's Law." *Philosophical Review* 96: 555–78.

Sainsbury, Mark. 1989. "What is a Vague Object?" *Analysis* 49: 99–103.

Salmon, Nathan. 1982. *Reference and Essence*. Oxford: Basil Blackwell.

Tye, Michael. 1996. "Fuzzy Realism and the Problem of the Many." *Philosophical Studies* 81: 215–25.

van Inwagen, Peter. 1988. "How to Reason with Vague Objects." *Philosophical Topics* 16: 255–84.

Wiggins, David. 1986. "On Singling Out an Object Determinately." In *Subject, Thought and Context*, 169–80, ed. P. Pettit and J. McDowell. Oxford: Oxford University Press.

Williamson, Timothy. 1994. *Vagueness*. London: Routledge.

————. Forthcoming. "Vagueness, Identity and Leibniz's Law." In *Individuals, Essence and Identity: Themes in Analytic Metaphysics*, ed. A. Bottami, M. Carrava and P. Giaretta, Dordrecht: Kluwer, Topoi Library, 2001 or 2002.

Shifting Sands: An Interest-Relative Theory of Vagueness

Delia Graff
Princeton University

I. INTRODUCTION

Saul Kripke pointed out that whether or not an utterance gives rise to a liar-like paradox cannot always be determined by checking just its form or content.[1] Whether or not Jones's utterance of 'Everything Nixon said is true' is paradoxical depends in part on what Nixon said. Something similar may be said about the sorites paradox. For example, whether or not the predicate 'are enough grains of coffee for Smith's purposes' gives rise to a sorites paradox depends at least in part on what Smith's purposes are. If Smith's purpose is to make some coffee to drink, so that he can wake up and start his day, then we would be inclined to accept, and would find it strange to deny, the following *sorites sentence:*

> For any *n*, if *n* grains of coffee are enough for Smith's purposes, then so are *n*–1.

Given his purpose, how could one grain make for the difference in the possibility of its being achieved? If, however, Smith has a stack of twenty quarters on one pan of a finely calibrated balance, and for whatever reason his purpose is to use the coffee grounds to tip the balance, then we would in no way be inclined to accept, or find it in the least strange to deny, the sorites sentence. If we have no idea what Smith's purposes are, then we should

have no attitude pro or con toward the sentence. We must conclude that whether or not a predicate is sorites susceptible does not depend only on the form or content of utterances of sentences containing it.

Do the preceding considerations threaten the view that all vagueness is vagueness in language? For it seems natural to say that what makes the sorites sentence paradoxical in one case but not in the other is that in one case but not the other Smith has a vague *purpose*. This should not be confused with the claim that 'purpose' is vague (though it surely is). Although it may be unclear at what point an act-type (e.g., *saving gorillas from extinction*) comes to satisfy the predicate 'is among Smith's purposes', we can just stipulate that, in the first case, *making some coffee* definitely is among his purposes, so that the paradox would arise even if 'purpose' were not vague. Let's say that doing such-and-such becomes one of Smith's *purposes** when and only when he has typed into his electronic diary 'I hereby intend to do such-and-such', and that once it is among his *purposes** it ceases to be when and only when he has typed into his diary 'I hereby no longer intend to do such-and-such'. Replace 'purposes' with the precise 'purposes*' in the above discussion, and the paradox does not go away.

Perhaps the vagueness of 'coffee', as applied to the beverage, is responsible for the paradox when Smith's purpose is as described in the first case. Start with a cup of coffee, keep replacing drops of it with equal amounts of spring water, stirring after each stage, and eventually the liquid in your cup will no longer be coffee, though it is difficult to accept that any one stage marks the transition. So it is true that 'coffee' is vague, but again, the paradox could arise even if it weren't. We can just stipulate that whatever number of grains of coffee Smith starts with, the beverage he makes with them will be a clear case of 'coffee'. (Or let 'coffee*' be a non-vague predicate, and suppose that Smith's purpose is to make some coffee* to drink, so that he may wake up and start his day.) Still, we would be inclined to accept that if n grains are enough for his purposes then so are $n-1$.

What about 'enough'; is it vague? I'm inclined to think not, or that if it is, it is only derivatively so. To say that n grains of coffee are *enough* for Smith's purposes is to say that were he to have n grains, he would not be prevented from achieving his purposes by not having more. Just as sentences have truth conditions and desires have satisfaction conditions, purposes have *achievement* conditions: those propositions, or states of affairs, which are such that if and only if one of them obtains will the purpose have been achieved. When we say that Smith's purpose is to make some coffee, typically we will not mean that there is some coffee, or even some amount of coffee, such that his purpose is to make *that* coffee, or *that* amount of coffee. (This is like when we say that Smith wants to catch a fish, we typically will not mean that there is a fish such that he wants to catch *it*.) But not just any amount of coffee will do. If he makes just a teaspoon, he will not

have achieved his purpose. We can imagine a *sorites series* of scenarios along which Smith makes increasingly greater amounts of coffee. We will find it hard to accept, if the increments are sufficiently small, that there is a first scenario in which it will be true to say that Smith's purpose has been achieved. Does this mean that 'achieve', and therefore 'enough', is vague? Let's approach this question by comparing 'achieve' as it applies to purposes with 'satisfies' as it applies to predicates.

It is often noted that if a predicate F is vague, then the predicate 'satisfies F' will also be vague. This leads some to say that 'satisfies' and also 'true' are vague because they "inherit" the vagueness of words like 'bald', 'heap', and 'tall'. The thought must be that if a complex expression such as 'satisfies 'tall'' is vague, then it must be because one of the words in the expression is vague. Since we may suppose that the expression ''tall'' occurring in 'satisfies 'tall'' is *not* vague, 'satisfies' must be vague, it being the only other word in the expression. If this little argument for thinking that 'satisfies' is vague is good, then we may still say that it is vague only in a derivative sense; its vagueness is parasitic on the vagueness of 'tall'. But why think that if a complex expression is vague it must be because one of the words in the expression is vague? It could be that containing a vague word is just one way for a complex expression to be vague, and that containing an expression that *denotes* something vague is another way for a complex expression to be vague. If this is correct, then since ''tall'' denotes a vague word, we have no reason to think that 'satisfies' is vague as well. The vagueness of the denotation of ''tall'' might just be sufficient for rendering the complex expression 'satisfies 'tall'' vague. Similarly, if the predicate 'is a scenario in which Smith's purpose has been achieved' is vague, it might be just because Smith's purpose (the denotation of 'Smith's purpose') is vague—that is, it is a purpose whose achievement conditions have a vague boundary—and if 'achieve' and hence 'enough' are vague too, their vagueness is parasitic on the vagueness of his purpose.

Purposes and desires can be vague because their achievement or satisfaction conditions may have vague boundaries. This could be true even if there were no vagueness in language (or at most only parasitic vagueness in words like 'achieve', 'satisfy', and 'true'). Pierre's desire for some champagne may be vague because of the vagueness in just how much champagne is required to satisfy it, even though the words we would normally use to describe that desire, 'some' and 'champagne', are perfectly precise, or at any rate, might as well be. (I gather the French have very strict and well-defined standards for what counts as 'champagne'.)

When I say that our purposes and desires can be vague, I do not mean that they have borderline cases—in the sense that there may be situations in which it seems correct to say that a purpose, for example, has not definitely been achieved nor has it definitely not been achieved. Many readers will

probably find this last remark strange, since philosophers regularly take it that the possibility of having borderline cases is *the* defining feature of vagueness. But a number of philosophers have come to doubt or reject this idea.[2] For if one thinks that a predicate has borderline cases just in case there's a gap between its extension and its anti-extension (or between its "definite" extension and its "definite" anti-extension), then it can seem that a predicate could have borderline cases even if the extent of that gap were perfectly precise, or even known. But then if having borderline cases were all there was to vagueness, there would be no obvious connection between vagueness and the sorites paradox. (Why would we be inclined to accept that any man 1 mm shorter than a tall man is tall if 'tall' were merely gappy?) Moreover, there seems to be a lack of fit between the borderline-case conception of vagueness and our metaphorical characterization of vagueness as a lack of sharp boundaries. The metaphor is traceable to Frege. In volume II of the *Grundgesetze* he writes:

> If we represent concepts in extension by areas on a plane, this is admittedly a picture that may be used only with caution, but here it can do us good service. To a concept without sharp boundary there would correspond an area that had not a sharp boundary-line all round, but in places just vaguely faded away into the background. (Geach and Black 1980: 139)

Frege goes on to say that this would not really be a concept at all. But if we take vagueness seriously, and if we recognize that we are only speaking metaphorically, the admission of *fuzzy* boundaries seems just what we want to characterize vagueness. The reason, then, for rejecting borderline cases as the defining mark of vagueness is that it seems that a predicate could have borderline cases without having fuzzy boundaries.

I suspect that many will want to affirm the converse: that if a predicate has fuzzy boundaries then it will have borderline cases. The thought may be this: a fuzzy boundary must occupy an extended region; to lie within this region is to be a borderline case. I worry, though, that this argument derives its entire force from taking Frege's metaphor too literally. In fact, I think we should be equally happy with a different metaphor, one that does not straightforwardly support such an argument. The metaphor, taken from R. M. Sainsbury, is one of *boundarylessness:* "vague concepts are concepts *without* boundaries," they "classify without setting boundaries" (Sainsbury 1991a: 6; emphasis added).[3]

I don't know whether there is any non-metaphorical yet theory-neutral way of characterizing vagueness. But we seem to have a clear enough grip on which phenomenon it is we're talking about for the metaphors to be useful. Sainsbury's metaphor, though, as compared with Frege's, has the advantage of more clearly bringing out the connection between vagueness and the sorites paradox. We're inclined to accept and equally disinclined to deny sorites sentences precisely because it seems to us that vague predicates

could draw *no* boundary, of any kind, between things or states of affairs that are sufficiently similar in the relevant respect. As Crispin Wright (1975) puts it, vague predicates seem to us to be always "tolerant" of small changes without always being tolerant of large ones. So when I say (to return finally) that our purposes and desires can be vague, I mean that we may have purposes and desires with achievement or satisfaction conditions that seem to us tolerant in the requisite sense, and hence boundaryless. When I desire some coffee, as I do every morning, how could it be that the desire I have could go from being satisfied to being unsatisfied—or even from being definitely satisfied to being borderline satisfied—just by removing one little drop of coffee from the cup I have in front of me?

The discussion so far has centered around the vagueness of our purposes and desires—let me call these our *interests*—but my main aim in this paper is to give an account of why vague *expressions* seem boundaryless to us. I'll begin by giving a semantic explanation of the phenomenon I want to account for, but will ultimately propose that the semantic explanation has psychological underpinnings. On the account I'll propose, the semantics of vague expressions renders the truth conditions of utterances containing them sensitive to our interests, with the result that vagueness in language has a traceable source in the vagueness of our interests. In the course of providing this account, I'll also explain why our interests can be vague, why it is they can seem tolerant—and hence boundaryless—in a way that leaves their coherence intact. Before proceeding in section III with my own proposal, however, I want to discuss some of the going solutions to the sorites paradox in order to focus ideas.

II. THREE QUESTIONS

Not all solutions to the paradox proposed by philosophers are designed to address just the same set of questions. In this section I want to set out some of the different questions the sorites paradox raises and to say which of those questions others have focused on and which of those questions will be my focus here.

I'll say that we have an instance of the sorites paradox when we are confronted with a group of sentences having the following form, each of which seems individually plausible:

(A) Fa (B) $(\forall x)(\forall y)(Fx \ \& \ Rxy \rightarrow Fy)$
(C) $\neg Fz$ (D) $(\exists b_1 \ldots b_n)(Rab_1 \ \& \ Rb_1b_2 \ \& \ldots \& \ Rb_{n-1}b_n \ \& \ Rb_nz)$

Here 'F' is to be a vague predicate; a is to be a clear case of the predicate; z is to be a clear non-case; and 'R' is to be replaced with some relation that renders sentences (B) and (D) plausible. So 'F' might be replaced with 'is a

tall man'; it might be that a is some professional basketball player and z is some professional jockey; in which case 'R' could stand in for the relation 'is 1 mm taller than', or maybe just 'is the predecessor of in s', where s is some appropriately constructed sorites series. Sentences with the form of (D) then say that *there is* a sorites series of the relevant kind with a as the first member and z as the last.

Since (appropriate instances of) (A–D) seem individually plausible but jointly inconsistent, a first question to address is: what has to give, one or more of (A–D), or their inconsistency? I follow most philosophers in giving up the truth of (B), the sorites sentence.[4] But as soon as one prepares to give up the truth of sorites sentences, a number of new problems arise:

1. *The Semantic Question.* If the universal generalization '$(\forall x)(\forall y)(Fx \& Rxy \rightarrow Fy)$' is not true, then must this classical equivalent of its negation be true?

   The "sharp boundaries" claim: $(\exists x)(\exists y)(Fx \& Rxy \& \neg Fy)$

   (a) If the sharp boundaries claim *is* true, how is its truth compatible with the fact that vague predicates have border-line cases? For the sharp boundaries claim seems to deny just that.
   (b) If the sharp boundaries claim is *not* true, then given that a classical equivalent of its negation is not true either, what revision of classical logic and semantics must be made to accommodate that fact?

2. *The Epistemological Question.* If '$(\forall x)(\forall y)(Fx \& Rxy \rightarrow Fy)$' is not true, why are we unable to say which one (or more) of its instances is not true—even when, say if the F in question is 'is a tall man', all the heights of the possible values of x and y are known?

3. *The Psychological Question.* If the universally generalized sorites sentence is not true, why were we so inclined to accept it in the first place? In other words, what is it about vague predicates that makes them seem tolerant, and hence boundaryless to us?

FINE'S SUPERVALUATIONISM

Let's first consider an account of the sorites based on a supervaluational treatment of vagueness, like that developed by Kit Fine (1975), according to which a sentence is true (false) just in case it is classically true (false) for all admissible ways of drawing precise boundaries for the vague expressions in the language. On Fine's view, borderline cases lead to truth-value gaps; that is, to be a borderline case of a predicate is to be in neither its extension nor its anti-extension. Although his supervaluation semantics renders sorites sentences false and sharp-boundaries claims true, supervaluationism still affords the following answer to question 1(a) above: Although sharp-boundaries claims are true, this is compatible with a predicate's having borderline

cases—construed as extension gaps—since sharp-boundaries claims for vague predicates, though true, have no true instances.

As it stands, however, Fine's account provides no answers to the Epistemological and Psychological Questions posed above: why are we unable to say which instances of a given sorites sentence are not true; and why were we so inclined to accept the sorites sentence in the first place? With regard to the Epistemological Question, on a truth-value gap approach to vagueness, sorites sentences will typically have a range of untrue instances. Suppose we have an appropriately constructed sorites series for a given vague predicate F, one on which each member of the series is R-related to the next. If we restrict the range of variables to just the members of such a series, then on Fine's account '$Fx \& Rxy \rightarrow Fy$' will have no truth-value just when either of the values of 'x' and 'y' is in the extension gap of F. (Whenever, that is, 'Rxy' is true). Thus to know which instances of the sorites sentence are untrue is just to know exactly which things on the series are in the extension gap of F. This is something we evidently do not know, and we have no explanation why. Fine's sophisticated treatment of higher-order vagueness does not in any way remedy this situation.[5]

With regard to the Psychological Question, Fine recognizes that there is at least some pressure to explain why, if sorites sentences for vague predicates are false, we are nevertheless so attracted to them. His remarks here are brief, however. I quote those remarks in full:

> I suspect that the temptation to say that [a sorites sentence] is true may have two causes. The first is that the value of a falsifying n appears to be arbitrary. This arbitrariness has nothing to do with vagueness as such. A similar case, but not involving vagueness, is: if n straws do not break a camel's back, then nor do $(n + 1)$ straws. The second cause is what one might call truth-value shift. This also lies behind LEM. Thus $A \vee \neg A$ holds in virtue of a truth that shifts from disjunct to disjunct for different complete specifications, just as the sentence 'for some n a man with n hairs is bald but a man with $n + 1$ hairs is not' is true for an n that shifts for different complete specifications. (Fine 1975: 286)

It is odd that Fine should cite the *arbitrariness* of the falsifying n as a cause of our temptation to say that a sorites sentence is true. In many cases, we readily acknowledge that a sentence with similar form is not true, or that a "slippery slope" argument is not a good one, despite the arbitrariness of the bad step. That is why the claim "That's the straw that breaks the camel's back" has such metaphorical weight. If I tug your ear once, you may not get angry. If I keep tugging your ear, when it is has long ceased to be funny, you will eventually get angry. There will be some arbitrariness in *which* of the tugs pushes you over the edge, but that in no way leads me to believe that if n ear tugs do not provoke your wrath, then neither will $n+1$.

The second cited cause of our temptation to accept sorites sentences is

a "truth-value shift." Here Fine must mean that we do not recognize sorites sentences as false because their falsity holds in virtue of a falsity that shifts from instance to instance for different complete specifications. One important worry here is that, especially for someone who rejects bivalence, an explanation of why we don't think a sentence is false should not count as an explanation of why we do think that sentence is true. A further but related worry is that the explanation seems tantamount to the claim that we are tempted to accept sorites sentences because we do not recognize supervaluational semantics to be correct. The claim raises a serious methodological question. The project of providing an adequate semantic theory of natural languages is an empirical enterprise. What are we to take as the data of our enterprise? Which are the phenomena against which we measure our theory for adequacy or correctness? What observed phenomena could prove a theory wrong? One of the main motivations for and attraction of the supervaluationist's theory is that it manages to verify sentences that express "penumbral connections"—such as 'If Jim is taller than Eric, then Jim is tall if Eric is'—where truth-functional theories that admit truth-value gaps do not. But why is our inclination to regard true existential and disjunctive claims as having a corresponding 'which one?' question that has a correct answer given any less weight? Fine seems to think that overall best fit with our intuitions is the best we can hope for in a semantic theory, even when sometimes the fit is not at all good at certain places.[6] Acceptance of bivalence would be one way to accommodate both the penumbral intuition as well as the clear intuition that true existential claims have a true instance. Fine does not consider the question whether acceptance of bivalence might yield an overall best fit.

When I say that Fine's account of vagueness provides an answer only to the first of the three questions I set out above, I do not mean to suggest that the account should for that reason be rejected. Rather, I want to point out that the account addresses some issues but not others, and to emphasize that without supplementation it cannot be regarded as a complete solution to the sorites.

WILLIAMSON'S EPISTEMICISM

Timothy Williamson is another philosopher who thinks that sorites sentences are false, and that sharp-boundaries claims are true. Unlike Fine, however, Williamson accepts bivalence, and so must answer the Semantic Question I posed by accounting for borderline cases in a different way from those philosophers who posit extension gaps for vague predicates. Williamson proposes instead that borderline cases of vague predicates are those things of which it's *unknowable* whether or not they're in the predicate's extension. Williamson aims to make this sort of ignorance plausible by arguing that just as we may have inexact knowledge of the number of

people in a crowded stadium, or of the location of what was the top card in a deck after the deck is cut, we may also have inexact knowledge of the meanings of expressions in our own language.[7] We know enough about the meaning of 'is a tall man' to know that any seven-foot tall man is in its extension; we don't know enough to know what the *least* height is that's sufficient for a man to be in its extension.

Williamson's account is well suited to provide an answer to the Epistemological Question I posed precisely because it gives an epistemological answer to the Semantic Question.[8] Given a bivalent semantics for vague predicates, sorites sentences will have exactly one untrue instance.[9] On Williamson's view, inexact knowledge explains why we don't know which instance that is. Still, no answer to the Psychological Question is forthcoming. Even if we accept Williamson's epistemicism, and so come to accept that sorites sentences are false, it still remains a mystery why we were so attracted to them in the first place.

In further developments of his view, Williamson does address the question why we can't even *imagine* discovering the precise location of the boundary for a vague predicate, whereas we could easily imagine discovering the precise number of people in a crowd, or the precise location of a certain card in a deck.[10] Given the sense of *unimaginability* that turns out to be at issue, however, these further developments have no bearing on the present question. The sense in which precise boundaries for vague predicates are argued to be unimaginable is this: when we imagine a sorites series for a vague predicate F, although there is a point in the imagined series that marks the transition from the Fs to the non-Fs, we cannot know where in the imagined series that point is. The source of our ignorance in the face of an imagined series is, as before, inexact knowledge of meaning. So, as before, although we do have an account of why we don't know where the boundary of F's extension is in an appropriately constructed series (whether imagined or not), we have no account of why we *believe*, of every particular point in the series, that the boundary is not there—why we believe of every point in the series that it does *not* mark the transition from the Fs to the non-Fs, or indeed from the definitely-Fs to things with any other status.

DEGREES OF TRUTH

I'll have little to say here about approaches to vagueness that adopt *degrees* of truth, but at least a few words are in order. One attraction of degree-theories, I would guess, is that it seems that they provide answers to both the Psychological and Epistemological Questions, while theories that admit of just two or three truth-values do not. Although the details may vary, any proponent of degrees of truth can say that we are inclined to accept sorites sentences as true because each instance of a sorites sentence has at worst a very high degree of truth; we are unable to say exactly which instances of a

sorites sentence are less than perfectly true because that would involve locating a boundary in a sorites series between a thing of which F is perfectly true and a thing of which it is only slightly less than perfectly true. The answers seem smooth enough. But in the absence of some substantial philosophical account of what degrees of truth are, we have no reason to accept that it should be both natural and common to mistake high degrees of truth for the highest degree of truth, or to mistake a small difference in degree of truth for no difference in degree of truth.

CONTEXT-DEPENDENT THEORIES

My purpose in discussing the theories of Fine and Williamson is to contrast them with theories that are designed specifically to address what I'm calling the Psychological Question. A growing number of linguists and philosophers appeal to the *context-dependence* of vague predicates in order to do just that.[11] Hans Kamp, for example, was prompted to reject the supervaluational semantics for vague predicates he developed in "Two Theories about Adjectives" (Kamp 1975) precisely because it remains mysterious on that account why vague predicates seem tolerant to us. On the semantic account of vagueness he later developed in "The Paradox of the Heap" (Kamp 1981), it turns out that vague predicates really are tolerant in the sense that every instance of a sorites sentence really is true. In other words, on the semantics Kamp came to favor, the sentence 'if x is tall and x is just 1 mm taller than y, then y is tall' is true, in any context, for any values of 'x' and 'y'.[12] In order to block unwanted conclusions, he proposes a radical revision of classical semantics, one that involves the notion of an *incoherent context,* and according to which a false universal generalization may have none but true instances, and on which a false conclusion may sometimes be inferred from true premises by only valid rules of inference. Though I think revisions this radical should be avoided if at all possible, I am in complete sympathy with Kamp's motivation. Since like Sainsbury I take apparent tolerance and boundarylessness— rather than borderline cases—to be the defining features of vagueness, I want an account that is geared to address the Epistemological and Psychological Questions I posed, since in effect, answering both of these questions requires us to explain *why* vague predicates seem tolerant to us, even though sorites reasoning shows us that they cannot be. I am happy to wait and see what story about borderline cases, and the characteristic hedging responses they provoke, might naturally flow from such an account.

III. THE BARE-BONES SOLUTION

My solution to the sorites is going to unfold in layers. The first layer is what I call the "bare-bones" account. It is most closely allied with those solutions to the paradox that appeal to the context-dependence of vague expressions,[13] especially those offered by Kamp (1981) and Soames (1999: chap. 7),[14] but is more neutral than these theories in a number of respects, perhaps most notably the question whether to accept bivalence. Discussion of another crucial respect in which the account is neutral will be postponed until the end of this section.

I'll begin by drawing attention to one commonly noted feature of vague expressions, namely that we can use them with *different standards* on different occasions. Sometimes the variation in standards can be traced to an implicit comparison class. For example, the sentence 'John is rich' might be uttered on one occasion to mean that John is rich *for a philosopher,* while it might be uttered on some different occasion to mean that John is rich *for an executive at Microsoft.* There has not been a great deal of attention paid in philosophical discussions of vagueness to the phenomenon of implicit comparison classes. The inattention is understandable, since making comparison classes explicit does nothing to resolve vagueness: the predicate 'tall for a basketball player' is no less vague than 'tall'. If the variation in the standards of use for vague expressions were always attributable to some variation in implicit comparison class, there would be little point in discussing that variation in the context of the present essay.

But it is not the case that the variation of the standards in use for a vague expression is always attributable to some implicit comparison class. Here is one example that lends some support to the claim. Suppose I am a casting agent auditioning actors for parts in a play. On one day I'm casting for someone to play the role of Yul Brynner, who had absolutely no hair. On a different day I'm casting for someone to play the role of Mikhail Gorbachev, who has some hair, but very little on top. When I turn away auditioners citing as a reason that they are not bald, I may be using different standards for 'bald' on the different days. I may say "Sorry, you're not bald" to an actor when he auditions to play Yul Brynner, and may then say, to that very same man when he auditions on the following day to play Gorbachev, "Yes you look the part; at least you're bald." My sense is that the variation in standard here is not due to a variation in implicit comparison class, since the comparison class implicitly at work in the two cases is the very same one—in each case, I meant 'bald for a man'.

Further argument for the claim that not every variation in standards is attributable to some variation in comparison class requires me to say a bit more about the workings of comparison classes. One point that's often either ignored or missed in discussions of the semantics of adjectives and their implicit or explicit relativization to comparison classes is that being tall for a basketball player, for example, is not just a matter of how your height compares with some *average* height for basketball players. If by some freak and tragic accident all the tall players are killed, so that the average height

of basketball players suddenly drops by a fairly large margin, it does not automatically become true to say of the tallest surviving player that he is tall for a basketball player. We can even lament the fact that none of the surviving players is tall for a basketball player.

Looked at another way, if by some freak and tragic accident it comes to be that all and only basketball players are golfers, it does not thereby become true to say that anyone who's tall for a golfer is tall for a basketball player. Whether one is tall for a certain kind depends on what the *typical* height is for things of that kind, and what the typical height is for a kind is not just some function of the heights of its presently existing instances.

What the preceding considerations are intended to show is that relativization to comparison classes is not an extensional phenomenon.[15] Comparison classes do not work just by contributing sets; for one, they need to form a kind. That is why it sounds strange to say that my computer is tall for a thing on my desk, even though it is in fact the tallest thing on my desk. Because the things on my desk don't form a kind, we have no notion of what a *typical* height is for a thing on my desk.[16]

My argument that not every variation in standards of use for vague expressions can be attributable to a variation in implicit comparison class can now proceed. I'll give an example of a variation in the standards of use for the adjective 'blue' where the only candidate comparison class does not form a kind. The example is a simple one. Suppose I want you to hand me a certain book. If the book in question is colored a very light grayish blue, and it's sitting among a bunch of other books, all of which are colored a very light grayish red, I may say, "Hand me the blue one." If, on the other hand, the book I want is sitting with a bunch of richly colored cobalt blue books, I may say, "Hand me the gray one." I take it that it would be true to say, in the first case, that the book I wanted was blue and, in the second case, that the book I wanted was gray. I also take it that 'gray' and 'blue' are mutually exclusive.

The variation in standards here obviously has something to do with the color of the books in the immediate surroundings of the one I want. But it cannot be that there is an implicit comparison class—one which would have been made explicit by using the expression 'blue for a book in its immediate surroundings'—because the books *in the immediate surroundings* do not form a kind, and so there is no notion of what a *typical* color is for things in that category. Also note that, in large measure, it is *up to me* whether I say 'gray' or 'blue'. I have some leeway and my choice is not completely dictated by the color of the other books lying around.

As a side point, let me add that the reason it sounds strange to say "That's blue for a book" is not that *books* (all of them) don't form a kind. They do. Rather, it sounds strange because, although books form a kind,

they come in such a wide array of colors, none of them more standard than the rest, that we have no notion of what a typical color is for a thing of that kind.

Returning to the main thread: we have some leeway in our standards of use for vague predicates; the variation in standards of use is not always attributable to a variation in comparison class; but still, we cannot use these predicates any old way we like.[17] What are the constraints? First, there are what we may call *clear-case* constraints. For each predicate, there will be only a limited range of cases which it will be permissible to count as positive instances. We can never use the word 'green' in such as way as to apply to the color of the sun. For each predicate there will also be a class of things which it will be mandatory to count as positive instances. No matter what standard is in place for 'blue', the predicate applies to the color of a clear afternoon sky. There will also be *relational* constraints for some predicates: whatever standard is in place for 'tall', anything the same height as or taller than something that meets the standard itself meets the standard. A further sort of constraint will *coordinate* the standards in use for related predicates: whatever standards are in use for 'rich' and 'poor', nothing can meet both, and it must be possible for something to meet neither. These three constraints together describe what the typical supervaluationist means by an "admissible" precisification of a vague predicate.

A fourth sort of constraint is what I call a *similarity* constraint. It is as follows: whatever standard is in use for a vague expression, anything that is *saliently* similar, in the relevant respect, to something that meets the standard itself meets the standard; anything saliently similar to something that fails to meet the standard itself fails to meet the standard. Put another way, if two things are saliently similar, then it cannot be that one is in the extension of a vague predicate, or in its anti-extension, while the other is not. If two things are similar in the relevant respect, but *not* saliently so, then it may be that one is in the extension, or in the anti-extension, of the predicate while the other is not. One reason for requiring that the similarity be a *salient* one is to block the absurd conclusions that would otherwise follow by sorites reasoning, since any two dissimilar things can be connected by a similarity chain.

The proposed constraints, then, fall into four categories: (i) Clear-Case Constraints, (ii) Relational Constraints, (iii) Coordination Constraints, and (iv) Similarity Constraints. Let me stress that these four types of constraint are being offered simply as *constraints*. It is not being suggested that they uniquely determine what standard of use for a vague expression must be in place in any given situation. I am merely claiming that the only standards of use that can be in place are ones that satisfy all the proposed constraints. It is also not being suggested that these constraints have any special *semantic*

status—that they are constraints which would be known, even implicitly, by any competent speaker. In fact, in the next section I'll deny that the Similarity Constraint is purely semantic, and will propose instead that it is in part a consequence of the vagueness of our interests.

I take it that the first three constraints I proposed are uncontroversial. But it is probably not obvious what the justification is for the fourth constraint. For the moment I want to provide justification for the fourth constraint by showing that it conforms neatly to our use. Suppose we're in an airport, and there are two suspicious-looking men I want to draw your attention to. I'm describing them so that you can pick them out of the crowd. But I can't point—that would be too conspicuous. You ask me, "Are they tall?" If the men are not much over five-foot, eleven inches, then (depending on the heights of the people in their immediate surroundings) I may have some leeway in choosing to answer "yes" or "no." But if, in addition, the two men are pretty much the same height—one is just noticeably shorter than the other—then the option is not available to me to say that one is tall but the other is not. Because the similarity of their heights is so perceptually salient—and now that you've asked me whether they're tall, also conversationally somewhat salient—I may not choose a standard that one meets but the other doesn't.

Another example, one that's slightly more far-fetched, is the following. Imagine an eccentric art collector who reserves one room for her paintings that contain just red pigments, and reserves another room for her paintings that contain just orange pigments. One day she is presented as a gift a painted color spectrum ranging from primary red on one end to orange on the other. She resolves to cut the canvas in half. Now if she cuts without thinking, perhaps in a state of mad excitement because she is so eccentric, she will most likely cut in just the right place—by which I mean that once the halves are re-framed and hung, she will still be able to truly proclaim that her paintings containing just red pigments are in one room, and that her paintings containing just orange pigments are in another. Although the right-hand edge of the painting in the red room is extremely similar to the lefthand edge of the painting in the orange room, their similarity is not *salient,* so the boundary between red and orange may occur between them. If the decision about where to cut is labored, in contrast, the collector will likely find herself unable to locate the boundary between the red and the orange, the pigments on either side of any proposed cut being too similar—and when the decision is labored, *saliently* similar—for one to go in the red room and the other in the orange.

I think we can come up with lots of examples of this kind.[18] A teacher might divide up her third-grade class into two groups, according to height. One group is to constitute the A-league basketball team, the other, the B-league team. Even if the heights of the children form a relatively smooth curve, it seems that she can truly say, once the division has become established, that the tall students are on the A-league team and the rest are on the B-league team. It would be true to say this given the division she's actually made, and it would have been true to say this even if, for auxiliary reasons, she had made the division at a slightly different point.

The preceding examples are intended to provide support for the Similarity Constraint by illustrating that when, but only when, one can manage to abstract from or ignore the extreme similarity between two objects— that is, when but only when their extreme similarity is not *salient*—the property one expresses in using a vague predicate is one which could be possessed by just one member of the pair. I concede that, despite the examples, the Similarity Constraint may still seem mysterious, even to those who at this point find it somewhat plausible. In the next section I'll give a fuller story about the notion of *salient similarity* adverted to in the Similarity Constraint, which I hope will have the double effect of making it seem both more plausible and less mysterious.

For now, however, I want to continue with discussion of the bare-bones account. In particular, I want to explain why, if the Similarity Constraint is descriptively correct, it provides the resources for answering the Epistemological and Psychological Questions. We have a universally generalized sorites sentence, for example the following:

$$(\forall x)(\forall y)(x \text{ is tall } \& \ y \text{ is just 1 mm shorter than } x \rightarrow y \text{ is tall})$$

In answer to the Epistemological Question, we may say that although it cannot be the case that every instance of such a universal generalization is true, the reason that we are unable to say just which instance or instances of it are not true is that when we evaluate any given instance, for any particular x and y that differ in height by just 1 mm, the very act of our evaluation raises the similarity of the pair to salience, which has the effect of rendering *true* the very instance we are considering. We cannot find the boundary of the extension of a vague predicate in a sorites series for that predicate, because the boundary can never be where we are looking. It shifts around. In answer to the Psychological Question, we may say that it is no wonder that we were so inclined in the first place to regard the universal generalization as true, given that any instance of it we consider is in fact true at the time we consider it.

The astute reader will no doubt have immediately noticed that my answers here presuppose bivalence in the sense that they do not follow from the Similarity Constraint *alone,* in absence of further semantic principles, unless bivalence is true. Let me explain. According to the Similarity Constraint for a given vague predicate F, when two things are not only very

similar in the respect relevant for predications of F, but also saliently so, then if one is in the extension of F so is the other, and if one is in the anti-extension of F so is the other. This leaves open the possibility, if there are truth-value gaps, that saliently similar things may both be in the extension gap of F, which in turn leaves open the possibility, in absence of further semantic principles, that an instance of a sorites sentence involving such a pair will itself be neither true nor false. My answers to the Epistemological and Psychological Questions required, in contrast, that such instances would not be merely valueless but true, as they would be given only bivalence and the Similarity Constraint.[19]

While it is true that I am an adherent of bivalence, it is not the case that the bare-bones solution is available only to those who accept bivalence. For example, a supervaluationist who is sympathetic to the account I've presented might hold (just as he does for the other three constraints I've discussed) that the Similarity Constraint acts as a constraint not only on the actual, partial extensions of vague predicates, but also on the precisifications of them that are operative in his truth-definition. To see how this would work, let us suppose that a and b are saliently similar in respect of height, and that a is the taller of the two. Given the Similarity Constraint, they are either both in the extension of 'tall', both in the anti-extension of 'tall', or both in its extension gap. Whichever one of these obtains, the conditional 'a is tall $\rightarrow$ b is tall' turns out to be supervaluationally true. For given the Similarity Constraint on precisifications, there will be no admissible way of drawing precise boundaries for 'tall' that divides the two. That is, on every admissible precisification of 'tall', either a and b are both in the extension of 'tall', or they are both out. Thus in every admissible precisification, the conditional 'a is tall $\rightarrow$ b is tall' is classically true, and hence the conditional is supervaluationally true. Thus the very answers to the Epistemological and Psychological Questions that follow from the Similarity Constraint, given acceptance of bivalence, are also available to the sympathetic supervaluationist.

I am not the first to propose that something like what I'm calling the Similarity Constraint is both descriptively correct and of use in solving the sorites. Kamp (1981) and Soames (1999), for example, also propose that the similarity of two objects can in special circumstances prevent the extension-boundaries of a vague predicate from occurring between them. Kamp and Soames differ from me, however, in respect of what those special circumstances are. The difference might be paraphrased as follows: while I propose that similar things cannot be divided by an extension-boundary for a vague predicate when their *similarity* is salient, Kamp and Soames propose that similar things can't be divided by an extension boundary for a vague predicate when it is salient that one of them is in the extension or that one of

them is in the anti-extension of the predicate. The repeated use of the word 'salient' here should not be given too much weight, however; it was merely intended as a useful way of bringing out the contrast. A more accurate way of expressing the similarity constraint these two adopt is as follows: if a and b are sufficiently similar and 'Fa' is "taken for granted by the participants in the discourse," then 'Fb' is true (Kamp); if a and b are sufficiently similar and a is "explicitly characterized" as falling under F, and "other conversational participants accept this," then 'Fb' is true (Soames).

A key difference between Kamp's and Soames's version on the one hand, and mine on the other, occurs in a situation in which it is part of the background of a conversation (in either Kamp's or Soames's sense) that a falls under F, but the similarity of a and b is *not* salient. In such a situation, given Kamp's and Soames's proposed constraint, b is in the extension of F, while given my proposed constraint it need not be. My proposal seems preferable in its handling of cases discussed above, like that of the eccentric art collector, or of the third-grade teacher who must divide her class into two basketball teams. Suppose that the collector is in her orange room, and is explaining to a guest that she hangs her paintings containing just orange pigments in that room, and that her paintings containing just red pigments are hung in another. Suppose also that they are examining the redder edge of the orange half of the spectrum, and that enough conversation has taken place so that it is now part of the background of the conversation, in whatever sense, that 'orange' is being used in such a way so as to apply to that redder edge. The similarity of that edge to the more orange edge of the other half in her red room is not at all salient. We may suppose that it has been years since the art collector has been in her red room or even thought about the fact that there is another half of the spectrum hanging there. The guest is not even aware of the existence of another half. Given my constraint, what the collector says to her guest can be true. Given Kamp's and Soames's constraint, it cannot.

Let's put this difference aside however. What's of greater interest here, given that Kamp and Soames each reject bivalence and that neither adopts supervaluation semantics, is how they put their own versions of the similarity constraint to use in solving the sorites paradox in a way that explains our initial temptation to regard sorites sentences as true. On this question they diverge. On Kamp's view, every instance of a sorites sentence is in fact true, in *any* context. The result is achieved by adopting a semantics for the conditional that is not truth-functional. A conditional is on this view true in a context c just in case either its antecedent is false in c, or its consequent *would* be true in the context that would result from c by incorporating the antecedent of the conditional into the background of c. This approach, combined with Kamp's version of the similarity constraint—whenever a and b

are similar and '*Fa*' is part of the background, '*Fb*' is true—has the effect of verifying, in every context, every instance of a sorites sentence.[20]

Soames, in contrast, adopts a straightforward truth-functional semantics for the connectives, that given by the strong Kleene truth tables, according to which a conditional with a valueless antecedent and consequent is itself valueless. Now if vague predicates have extension gaps, both Soames's version and my version of the similarity constraint allow for the possibility that similar objects may both be in the extension gap of a vague predicate, even if their similarity is salient. Thus if Adam is just 1 mm taller than Bert, and this extreme similarity between their heights is not only known but salient, it may still be, given the strong Kleene tables, that the sentence 'If Adam is tall, then Bert is tall' has no truth value. The answer I offered to the Psychological Question is therefore not available to Soames.[21] For I had proposed that our inclination to accept sorites sentences and our inability to say just which instance or instances of them are not true is explained by the fact that any instance we consider is, as we consider it, in fact true. Soames, also concerned to answer the Psychological Question, proposes instead that the explanation for our initial attraction to sorites sentences is that we *mistake* them for true *meta-linguistic* principles, namely, those principles that give expression to his version of the similarity constraint.[22]

Soames's explanation requires for its force that the truth of these meta-linguistic principles is something widely known. I gather that Soames is thinking of his version of the similarity constraint as being a feature of the *meaning* of a vague predicate, and that therefore we would know, at least implicitly, that similarity constraints are true in virtue of our being competent speakers. Whether similarity constraints have this sort of semantic status is something which should be open to question, however. As I further develop my own view in the next section, I'll suggest that this is precisely not the case—that similarity constraints are empirical truths, made true, at least in part, because we have the kinds of interests that we do.

My reason for discussing the views of Kamp and Soames here is not merely to explain how my own view compares with others related to it, but also, in so doing, to emphasize that something very much like, if not exactly like, the bare-bones solution to the sorites is available to those with a variety of semantical and logical commitments and that, in particular, it is neutral with respect to the question of bivalence.

I should mention two other philosophers whom I think of as being in the bare-bones camp, even though their views go well beyond the skeletal account offered here. Unfortunately, my remarks must be briefer than I would like. Diana Raffman employs something like a similarity constraint for vague predicates in her solution to the sorites paradox. On her view, there is a true object-language principle, which she calls 'IP*', expressed as

follows for the specific case of the predicate 'looks red': "for any *n*, if patch #*n* looks red then patch #(*n*+1) looks red, *insofar as #n and #(n+1) are judged pairwise*" (Raffman 1994: 47). On Raffman's view, as on Soames's our inclination to accept sorites sentences is explained by a proposed tendency to mistake them for distinct but related claims—in Raffman's case, for the true claim IP*.[23] In one way, my own Similarity Constraint seems more closely related to Raffman's version than to either of Kamp's or Soames's, in that it seems that on Raffman's view the *salience* of a given similarity (in her case, perceptual salience) is doing the crucial work.

Jamie Tappenden's (1993) view is a more distant relative, but still, something like a similarity constraint plays a central role in his account of the sorites. Tappenden, who also admits truth-value gaps, proposes that if two objects are sufficiently similar (whether or not that similarity is salient) then the standards in use for a vague predicate *F* must not be such as to place one of the objects in the extension of *F* and the other in its anti-extension. (Extension/Gap and Gap/Anti-Extension pairs of assignments are both fine.) Given also his acceptance of the strong Kleene tables, the inadmissibility of using such standards amounts to the inadmissibility of using a vague predicate in such a way as to render a sorites sentence for it *false*. This privileged status, *inadmissible to falsify*, is argued to support a special kind of speech act which Tappenden calls "articulation." The point of articulating a sentence is to get someone to stop or refrain from using a vague predicate in an inadmissible way, where one way to use a vague predicate inadmissibly is to use it in a way that does not conform to Tappenden's version of the similarity constraint. Tappenden explains our inclination to accept sorites sentences by claiming that since, though untrue, they are articulable, we may be warranted in uttering them; we mistake our good reason for uttering sorites sentences as reason to believe them.[24] One problem for the account is that sharp-boundaries claims (also not true according to Tappenden) turn out to have the same privileged status as sorites sentences. Something other than *articulability* is required to explain the difference in our attitudes toward such sentences.

I want to conclude this section with some remarks about the respects in which the bare-bones account is bare-boned. I have emphasized that the account is neutral on the question of bivalence. It has also been silent about how exactly to cash out the notion of *salience* that has been used so freely. We should probably think of the occurrences of 'salient' and its cognates as they've occurred here as just a peg on which to hang some more substantive theory, which I aim to provide. The account presented so far is also bare-boned in another respect: it is neutral about how the variation in standards of use for vague predicates should be characterized semantically. In particular, it has been left open just how much of this variation should be thought

of as arising from *context-dependence*. Let me explain. When I say that the standards of use for a predicate like 'tall' can vary, I mean that the extension of 'tall' can change from occasion to occasion, *even if the heights of everything remain stable*. When I say that the standards for 'expensive' can vary, I mean that the extension of the predicate can vary, *even if the costs of everything remain stable*. Given that the extensions of vague predicates can vary in this way, we might want to say that every such variation points to a change in context. One idea might be that 'tall' always expresses an intrinsic property, one which whether a thing has it depends only on that thing's height. Another idea might be that 'tall' always expresses a relational property, one which whether a thing has it depends only on how that thing's height compares to the height of certain other things. Either way, if the extension of 'tall' can vary even as the heights of *everything* remain stable, this could only be because 'tall' was context-dependent—because it expressed different properties on different occasions.

We should not assume off the bat, however, that 'tall' does express a property possession of which depends *only* on heights. It is possible that the predicate could express the same property from occasion to occasion, and that the reason the extension may change as the heights of things do not change is that the property expressed context-invariantly by 'tall' is a property which is such that whether a thing has it depends not only on heights, but on other things as well. I will go on to propose that despite the constant changing of standards of use for vague predicates that is dictated by the Similarity Constraint, there is much less context-dependence than one might have initially thought would be a consequence of the bare-bones view.

IV. VAGUE EXPRESSIONS ARE INTEREST-RELATIVE

The second layer of my solution to the sorites rests on a certain view about what predications involving vague expressions mean. Let me illustrate the view by means of examples. I propose that

That car is expensive

is to be analyzed as meaning:

That car costs a lot,

which in turn is to be analyzed as meaning:

That car costs significantly more than is typical.

Similarly, 'John is tall' is to be analyzed as meaning 'John has a lot of height', which is in turn to be analyzed as meaning 'John has significantly more height than is typical'. This sort of analysis is extended also to cases involving explicit relativization to a comparison class as follows: 'Mickey is old for a mouse' is to be analyzed as meaning 'Mickey has a lot of age for a mouse', which is in turn to be analyzed as meaning 'Mickey has significantly more age than is typical for a mouse'.

I do not want to try to get too specific about exactly what sense of 'analysis' is being assumed here. Suffice it to say, for the time being, that when I say that 'John is tall' is to be *analyzed* as meaning that John has significantly more height than is typical, I intend to be making a claim about what type of property it is that we attribute to John when we say that he is tall. In particular, I intend to be claiming that the property thereby attributed to John is not an intrinsic property, but rather a relational one. Moreover, it is not a property the possession of which depends only on the difference between John's height and some norm, but also on whether that difference is a significant one. I take it that whether or not a difference is a significant difference does not depend only on its magnitude, but also on what our interests are. (Exactly whose interests are at issue is a delicate question, which I'll try to ignore as much as possible.) The central claim of this section, indeed of the paper, is that if this interest-relative analysis is correct, then the Similarity Constraint, and the bare-bones solution along with it, drops out as a consequence.

I should first say a few words about my analysis of 'a lot' as meaning 'significantly more than is typical'. There is actually quite a bit more complexity to the issue. Sometimes 'a lot' can mean 'significantly more than is typical', but other times it can mean 'significantly more than is wanted or needed'. And also it can mean 'significantly more than is expected'. I feel sure that even these options are not exhaustive. It seems to me we need the 'typical' meaning for a case like the following. I am throwing a huge party. I know that there are going to be at least one hundred people coming. I have already done some shopping for the party, in particular, I have already stocked my refrigerator with beer. My friend Linda comes over to help with the preparations, and upon opening my refrigerator proclaims, "Wow! That's a lot of beer." Now, I have a normal size refrigerator—there are at most one hundred bottles in there. This is not enough beer for the huge party I'm throwing, and Linda knows this. Still, it is appropriate for her to say that there is *a lot* of beer in the refrigerator, because it is significantly more beer than one typically finds in a refrigerator.

The reason we cannot end the story there, however, is that analyzing 'a lot' as meaning 'significantly more than is typical' is not going to work for other sorts of cases. For example, it is both typical and expected that supermarkets will always have *a lot* of milk in stock. If 'a lot' just meant 'significantly more than is typical', then we could not say of our local supermarket

that it has a lot of milk in stock, since the amount of milk it has in stock is a completely typical amount for supermarkets to have. Here I think that when we say that the supermarket keeps a lot of milk in stock, we mean something like: it stocks significantly more milk than will be wanted or needed by shoppers on any given day. We cannot extend this proposal back to the beer case, however, since as I said, it is going to be a huge party, and I do not have *enough* beer for it.

What about a case like this: we're driving through Iowa, and suddenly we come upon a huge field of corn. We can see that the field stretches for miles. I say, "Wow, that's a lot of corn." Now it is not the case that I mean here that that's significantly more corn than is typical. We are in Iowa after all. Nor do I mean that it's significantly more than is wanted or needed. It's a big country, and we need a lot of corn. My exclamation, "Wow, that's a lot of corn" is my way of expressing some surprise—there is significantly more corn than I expected to come upon at that moment.

I want to stress that, despite these considerations, I am not wedded to the idea that 'a lot' is ambiguous. Perhaps it could be argued that the 'typical' meaning for 'a lot' really is sufficient, given some pragmatic considerations, to handle cases of the types I've discussed. The crucial feature of what I've said for my purposes here, is that 'a lot' means '*significantly* more than some norm'. Clearly, there is more than one kind of norm. It seems equally clear that the different uses of 'a lot' discussed above should in some way or another be accounted for as resulting from a difference in the kind of norm involved. Whether 'a lot' should therefore be deemed either ambiguous or context-dependent is a highly theoretical and largely theory-internal question which it seems to me could be plausibly answered either way.

Nevertheless, if vague expressions are to be analyzed as I've proposed, then we should expect that, like 'a lot', they could be used in ways that involve different norms on different occasions. Consider the expression 'old for a dog'. Here we have a vague adjective explicitly relativized to a comparison class. On my proposal, the expression is to be analyzed as meaning 'has significantly more age than is typical for a dog'. (Henceforth, 'typical' is to be understood as making reference to some or other norm.) Suppose that Fido is fourteen years old and Rover is twenty years old. Someone who says that Rover is old for a dog may be making a remark about his extreme longevity, while someone who says that Fido is old for a dog may be merely remarking that he is in his old age. This is to be explained, on my view, by a variation in norm. Rover, the twenty-year-old, has significantly more age than is the norm for a dog to attain; while Fido, the fourteen-year-old, has significantly more age than some different kind of norm for a dog, one that's much harder to articulate, but which perhaps concerns the peak age of good health.

There are also norms of expectation. This is what explains why it can be appropriate for me to say, when I see my young nephew for the first time in months at a family gathering, "Derek, you're so *tall*." It can be appropriate for me to say this even though I know that my nephew has always been and still is short for his age. What I am saying is that he has significantly more height than I expected him to have, given what his height was the last time I saw him.

Let us now shift our attention away from that feature of my proposal which has it that predications involving vague expressions involve reference to some norm, and turn our attention to the feature I am most concerned with, namely, that in claiming that someone is tall for a ten-year-old, for example, we are claiming that he has *significantly* more height than some norm for ten-year-olds. I want now to proceed with my central claim that if this is correct, the Similarity Constraint drops out as a consequence. The argument rests on the idea that two things that are qualitatively different in some respect, even when they are known to be different, can nonetheless be the *same for present purposes*. What I claim is that if two things are the same for present purposes, in respect of height say, then one can have *significantly* more height than is typical if and only if the other does. Given my proposed analysis of 'tall', it follows that when two things are in respect of height the same for present purposes, one is in the extension of 'tall' if and only if the other is. What's required, if the Similarity Constraint is to follow, is that two things come to be the same for present purposes when they are not only very similar, but when also their similarity is in some sense salient—in particular, when it is being actively considered.

Often it is appropriate to say of distinct but very similar things that they are the same for present purposes. For example, suppose a small child is watching me make a pot of coffee and, thinking she is being helpful, points out that a couple of grains have spilled from my coffee scoop. She assumes that the coffee scoop is a measuring device rather than a tool for convenient transfer, and so mistakenly concludes that this is an exact science. When she then wonders why I don't bother to replace the grounds I've spilled, I might explain to her that there is no need because the two amounts are the same for present purposes. Assuming that it is indeed *true* that the two amounts are the same for present purposes, this cannot entail that for present purposes: it is true that the two amounts are the same. It is false that the two amounts are the same, and my purposes cannot alter that fact. To say that the two amounts of coffee are the same for present purposes is to say that as far as my coffee-making purpose is concerned, the two amounts might as well be the same; my coffee-making purpose permits me to behave as if the two amounts were the same, since the purpose is in no way thwarted by my behaving as if they were the same.

Although my coffee-making purpose permits me to behave as if the two amounts are the same, what really makes the two amounts the same for present purposes is that my interests on the whole *require* me to behave as if the two amounts were the same. The reason for the requirement is that in addition to having an interest in getting a decent amount of coffee made, I also have in interest in doing this efficiently. Because of my interest in efficiency, there is some cost associated with measuring out the coffee grains precisely, and also some cost associated with taking the time to replace a few spilled grains. The cost of discriminating between the two amounts of coffee at hand clearly outweighs the benefits of discriminating between them, since even though I always prefer greater caffeine intake to less, I would be unlikely to notice, much less care a lot about, the differential effect if any these two amounts would have on me. From this point on, I will say that two things are the same (in a certain respect) for present purposes when the cost of discriminating between them (in that respect) outweighs the benefits. In saying this, I am not making a claim about what we *mean* when we say that two things are the same for present purposes. I think we just mean by this that for present purposes, it is fine to ignore the difference between the two things. Rather, I am claiming that what makes it fine to ignore the difference, is that the cost of discriminating outweighs the benefits.[25]

Crucially, there can only be a cost to discriminating between two things if they are in some sense both "live options." When I make my coffee in the morning and spill a few grains from the scoop, there is a cost to discriminating between the two amounts at hand. That is, there is a cost to taking the time to count and replace the grains I've spilled. The cost outweighs the benefits, so the two amounts are the same for present purposes. It does not follow, however, that any pair of amounts of coffee grains that differ to the same degree are also the same for present purposes. If it is not a live option that I'll use either of some pair of amounts, then there is no cost associated with discriminating between them, and hence they are *not* the same for present purposes.

Being the same for present purposes can be seen to play exactly the same role here as *being saliently similar* was playing in the bare-bones account. Thus I am cashing out the notions of "similarity" and "salient similarity" in the following way: to say that two things are similar enough to be subject to the Similarity Constraint is to say that were they live options, the cost of discriminating between them would outweigh the benefits. To say that they are *saliently* similar is to say that they are indeed live options, that the cost of discriminating between them *does* outweigh the benefits—that is, that they are in fact the same for present purposes. Given my analysis of 'tall', for example, as meaning 'has *significantly* more height than is typical', and given also that if two things are in respect of height the same for

present purposes one can have significantly more height than is typical if and only if the other does, the Similarity Constraint for 'tall' drops out as a consequence.

One potential problem that immediately arises is that if it can be true that two people are in respect of height the same for present purposes *whenever* their heights are at most 1 mm apart, that is, if there can be enough live options, then we could have it that whenever the heights of two people are at most 1 mm apart, one is in the extension of 'tall' if and only if the other is. (I'll just assume, for the sake of the objection, that the distribution of heights forms a sufficiently smooth curve.) Given the transitivity of the biconditional, this would have precisely the consequence we're trying to avoid, namely, that Michael Jordan is in the extension of 'tall' if and only if I am. Yet I am clearly not tall. To avoid this result, we need at least one pair of people whose heights are at most 1 mm apart yet whose heights are not the same for present purposes.

This is in fact a worry I should have mentioned when presenting the bare-bones version of my solution. The worry stated as it would have arisen there is that if we could get ourselves into a situation where the similarity in height of any two people that differed by at most 1 mm was a salient similarity, then by my Similarity Constraint it would follow that given any two people at all, one would be in the extension of 'tall' if and only if the other was. If every similarity in height is salient, then there is no place for the boundary between the tall and the not-tall to be. But this seems to be exactly the situation we are in when we are discussing the sorites paradox, and imagining that we are confronted with a series of people each of whom is just 1 mm taller than the next. The worry here is a serious one, for given also the Clear-Case Constraint, it may seem that my account is not merely false, but contradictory. For I am a clear case of 'not-tall'. So by the Clear-Case Constraint I am in the anti-extension of 'tall'. Given enough salient similarities, it would follow by the Similarity Constraint that Michael Jordan is in the anti-extension of 'tall', and hence not tall. But Michael Jordan is a clear-case of 'tall', and so by the Clear-Case Constraint, he is in the extension of 'tall', and therefore tall. Contradiction.

One way around this problem would be to argue that there cannot be enough salient similarities. When we are confronted with a sorites series for 'tall', although it is salient that each adjacent pair in the series is very similar, it is not the case that each adjacent pair in the series is such that it is salient that they are similar. There are too many pairs for us to actively entertain each similarity.

Another way around the problem we are facing would proceed by noting that the presented *reductio* of the bare-bones account contained a crucial mistake. It assumed that on any given occasion, there will be a standard in

use for a vague predicate that satisfies all of my constraints. But that is not quite right. Rather, it is that the only standard that can be in use for a vague predicate is one that satisfies all four constraints. In a situation where there are too many salient similarities, no standard could satisfy all four constraints. There is no contradiction in that. Thus the problem of too many salient similarities might be avoided by saying that in such a situation, the vague predicate at issue fails to express any property at all. In such a situation, no proposition is expressed by an utterance of a sentence containing the given vague predicate.

I actually think that once *being saliently similar* is cashed out as I've described, the first way around the problem proves to be the correct one. Whichever of the wide variety of purposes we may have in conversations where we have occasion to use vague expressions like 'tall', 'old', 'expensive', and the like, it will simply be a brute fact that that there will be a least height, age, or cost of which it is true to say that it is significantly greater than is typical—at least, that is, if it is true to say of any height, age, or cost at all that it is significantly greater than is typical. Any lesser amount simply cannot be the same for whatever purposes are in place. Even if there is some cost associated with any discrimination we might make between similar heights, ages, etc., the cost of making the discrimination nowhere does not outweigh the benefits of making it somewhere. The boundary between those differences that are significant and those that are not will try to locate itself, so to speak, at a place where there is least resistance.

I take it, returning to the example we began with, that when I desire some coffee it is just a brute fact that there is a least amount of coffee of which it is true to say that it will satisfy my desire. Anything less will not do. I know that many philosophers will protest; they will say, "But how could it be that your desire for coffee is like *that*?" I say, given that a teaspoon of coffee is not enough to satisfy my desire, how could it *not* be like that! Moreover, I have an explanation for why my desire seems that it is not like that. There should be nothing surprising or doubtful in the suggestion that we have inexact knowledge (in Williamson's sense) of the satisfaction conditions of our desires. I have never bothered to figure out exactly how the enjoyment I get from coffee maps on to the different amounts of it I might drink. And even were I to try to ascertain this, it would be close to impossible to ensure a controlled experiment, since it will always be that factors other than the amount I drink affect my enjoyment. But it is not merely that I have inexact knowledge of the satisfaction conditions of my desire for coffee on a given occasion; it is also that the satisfaction conditions of my desire may subtly shift, so as to be satisfied by different amounts of coffee as different options become available to me and the costs of discriminating between different pairs of amounts change. That, I claim, is the

essence of a vague desire, and it explains why we are inclined to accept sorites sentences such as the one that opened the essay, namely, 'When my purpose is just to make some coffee to drink, then for any n, if n grains of coffee are enough for my purposes, then so are $n-1$'.

V. A FORMAL PROPOSAL FOR VAGUE ADJECTIVES

I have not offered a general recipe for making the interest-relativity of vague expressions explicit. I have merely illustrated my view by discussion of a small number of examples, all of which have been vague adjectives. Not only have all of the vague expressions discussed been adjectives, they have also all been positive adjectives—positive in the sense that anything having *more* height, age, or cost than a tall, old, or expensive thing is itself tall, old, or expensive. What about the negative counterparts of these expressions? Well, to be short is to have significantly *less* height than is typical, or perhaps to *lack* significantly more height than is typical. To be young is to lack significantly more age than is typical. To be cheap is to lack significantly more cost than is typical.

It is not obvious, however, how we are to extend such patterns to vague expressions in other categories; for example, to vague nouns like 'heap'. To be a heap of Fs is to be a collection of Fs that is in an arrangement that is significantly like . . . , well, like what? I have some idea of what an arrangement of Fs must be like in order for it to be a heap of Fs. It has to be an arrangement significantly like one which could be the result of a process of pouring some Fs from a stable point above a surface. That is why pyramids of soup cans in the supermarkets are not heaps. Moreover, the pouring must continue for long enough for the height of the arrangement to be at least half the width of the base of the arrangement. That is why (some) mounds are not heaps.

There is a good explanation for why the treatment of vague adjectives runs smoothly, while extension of the proposal to vague nouns seems strained. The reason is that it is a semantic feature of adjectives that they are associated with some dimension of variation—one needed for the formation of comparatives such as 'taller', 'older', and 'more expensive'. But it is not a semantic feature of nouns that they are associated with a dimension of variation. That is why a generalization of my proposal to nouns would require a case-by-case analysis. I won't attempt any such generalization here. Rather, I will show how my proposal, confined to vague adjectives, might be formally implemented in a compositional semantics.

Throughout this essay, I have been speaking of adjectives such as 'tall' and 'old' as *predicates,* and I have been talking about their *extensions.* There

has been significant debate, however, about whether the thought that adjectives are predicates can be maintained. A major issue to contend with is that adjectives may occur in a wide variety of constructions. They may occur on their own in predicative position, as in 'John is tall'; they may occur in predicative position with modifiers such as 'very' and 'fairly' or with a prepositional phrase representing a comparison class, as in 'John is thin for a wrestler'; they may occur in prenominal (also called 'attributive') position, as in 'Fido is a big Chihuahua' and 'The little Chihuahua is mine'; they may occur in comparative and equative constructions, as in 'Fido is younger than Rover' and 'Peter is as rich as Paul'; and they may occur with 'too' or 'enough' and infinitival complements, as in 'Tom is too tall to be my skating partner' and 'Timmy is tall enough to go on this roller coaster'. One question is whether predicates are of the right semantic type to be accommodated compositionally in this wide variety of constructions.[26] A further but related question is how the meaning of such a complex expression containing an adjective could be a function of the meaning of the adjective occurring in it if adjectives were just predicates. For example, even if there is no type mismatch in combining an adjective such as 'young' with the comparative morpheme '-er' to form a relation expression, it is difficult to see how the meaning of the resulting relation expression could possibly be a function of the meaning of 'young' if it were just a predicate. Whether or not Fido is younger than Rover really has nothing to do with whether either is young. As Kamp (1975) puts it: "It is quite obvious that if adjectives were ordinary predicates no such transformation could exist. How could we possibly define the relation x *is bigger than* y in terms of nothing more than the extension of the alleged predicate *big*?" (127).

Nevertheless, Kamp (1975) does defend a predicate interpretation for adjectives (as does Ewan Klein (1980)). Montague (1970) and Cresswell (1976), in contrast, take the prenominal occurrence of adjectives to be primary and propose that adjectives are essentially noun modifiers: for Montague, the semantic function of an adjective is to combine with a noun to form a predicate; for Cresswell, adjectives combine with nouns to form relations that hold between individuals and degrees on a scale. Others (e.g., von Stechow (1984)) take adjectives themselves to stand for relations that hold between individuals and degrees.

The semantic account of adjectives I find most attractive, developed and defended by Christopher Kennedy (1999), constitutes yet another alternative. I won't undertake an explanation here of why I prefer the account to existing alternatives, or why I think it requires improvement. I'll only provide just the barest sketch of the account, and explain how it can be adapted to accommodate my own interest-relative proposal. The crux of the account is that adjectives do not have predicate-type semantic values, but instead denote functions from individuals to degrees on a scale. As I have indicated at a number of points, I am a strict adherent of bivalence, so these degrees must not be thought of as degrees of truth. 'Tall' denotes a function from individuals in the domain to degrees on the height scale—the taller the individual, the greater the degree. 'Short' also denotes a function from individuals to degrees on the height scale, but the degrees in the range of this function are of a different sort than those in the range of the function denoted by 'tall'. While 'tall' denotes a function from individuals to *positive* degrees on the height scale, 'short' denotes a function from individuals to *negative* degrees on the height scale. The shorter the individual, the greater (roughly speaking) is the negative degree assigned to it. Intuitively, a positive degree of height is the amount of height an individual has, while a negative degree of height is the amount of height it lacks. 'Expensive' denotes a function from individuals to degrees on the cost scale. 'Hot' denotes a function from individuals to degrees on the temperature scale, and so forth. Following Kennedy, I will call such functions 'measure functions'.

Before I explain how predicative uses of adjectives can be handled by such a view, it will be instructive to explain how the view deals with comparative constructions. Take a sentence containing a comparative construction like:

John is taller than Mary.

Here 'tall' as it occurs in 'tall-er' denotes a measure function. The semantic value of the degree morpheme '-er' is such as to combine with a measure function to yield a relation that holds between individuals. (Here the semantic value of a relation-expression is taken to be a function which given an individual as argument yields as value a function from individuals to truth values.) In particular the meaning of the degree morpheme '-er' can be represented as follows, using lambda abstraction notation to represent functions, as is familiar from Montague (1973):

$$-er = \lambda G \lambda y \lambda x (G(x) > G(y)).[27]$$

Here 'x' and 'y' are individual-level variables. And 'G' is to be a variable that ranges over measure functions. ('G' stands for 'gradable adjective'.) Thus for the semantic value of 'taller' we get:

$$taller = -er(tall) = \lambda y \lambda x (tall(x) > tall(y)).$$

Taking 'than' to be semantically vacuous we get the following as the semantic value of 'taller than Mary':

$$taller\ than\ Mary = taller(Mary) = \lambda x (tall(x) > tall(Mary)).$$

And then the following as the semantic value of the sentence:

John is taller than Mary = (tall(John) > tall(Mary)),

which equals the value *true* just in case the measure function denoted by 'tall' assigns a greater degree of height to John than it does to Mary—that is, just in case John's height is greater than Mary's, which is what we want.

So far so good. In cases where we have a predicative occurrence of an adjective with an explicit comparison class, as in

John is tall for a basketball player,

it is proposed that there is actually a phonologically null degree morpheme, represented here by '$\varnothing_{ABS}$' which Kennedy calls the *absolute* morpheme.[28] So just as we have 'John is tall-er than Mary', we have 'John is tall-$\varnothing_{ABS}$ for a basketball player'. The semantic value of the unpronounced absolute morpheme is such as to combine with a measure function to yield a relation that holds between individuals and properties. Kennedy proposes that while the comparative morpheme '-er' has roughly a *strictly greater than* meaning, the absolute morpheme has roughly an *at least as great as* meaning. I'll modify Kennedy's proposal in order to deal with the sorites paradox in the way I've outlined by saying that the absolute morpheme has roughly a *significantly greater than* meaning. More accurately, the meaning of the absolute morpheme $\varnothing_{ABS}$ I'll propose, is as follows:

$$\varnothing_{ABS} = \lambda G \lambda P \lambda x (G(x) \; !> (NORM(G))(P)).$$

Here 'G' and 'x' are as before, and 'P' is a variable that ranges over properties (or perhaps kinds). The greater than sign with an exclamation point before it is to stand for the *significantly greater than* relation. (The difference between my proposal and Kennedy's is that he has '$\geq$' where I have '$!>$'.) NORM combines with a measure function to yield a function from properties to degrees on the scale associated with the measure function. Intuitively, (NORM(G))(P) is the norm, or typical, amount of G-ness for things with property P. (NORM(*tall*))(the property of being a basketball player) is the norm or typical height for basketball players. Thus the absolute degree morpheme combines with 'tall' to yield the following relation between individuals and properties:

$$\varnothing_{ABS}(tall) = \lambda P \lambda x (tall(x) \; !> (NORM(tall))(P)).$$

'Tall for a basketball player' then has the following function as its semantic value:

($\varnothing_{ABS}(tall)$)(the property of being a basketball player) =
$\lambda x (tall(x) \; !> (NORM(tall))$(the property of being a basketball player)),

which assigns *true* to an individual x just in case x's height is significantly greater than is typical for a basketball player—as proposed. In cases where

there is no explicit comparison class, I assume that there is a comparison-class variable at logical form that gets assigned a property as value by the context.

Now that we have (a sketch of) a concrete proposal in hand, I should make some brief remarks about my earlier comment that on my account there is less context-dependence than one would have initially thought would be a consequence of the bare-bones view. Predications such as 'John is tall' are certainly context-dependent on the view I propose. One way in which the content of an utterance of this sentence depends on context stems from the need for implicit relativization to a comparison class. Other sources of context-dependence are the two context-dependent elements appealed to in stating the meaning of the unpronounced absolute morpheme $\varnothing_{ABS}$—namely, '!>' and 'NORM'. 'NORM' is context-dependent since, as was discussed in the last section, there are many different kinds of norms. *Significantly greater than* is a context-dependent relation, since what is significant to one person may not be significant to another. Any use of 'significant', or of any word whose content involves what is significant, requires an implicit subject with interests—an answer to the question: *significant to whom*? In saying that there is less context-dependence than one would have initially thought would follow from the bare-bones view, what I had in mind was this: the property attributed to John by a particular utterance of 'John is tall'—that is, once all contextual elements are fixed—is still a property the extension of which may vary even as the heights of everything remain stable, since the extension of the property may vary as the interests of the relevant parties vary, that is, as different differences become more and less significant as different similarities become more and less salient.

VI. THE SEMANTIC QUESTION

As an adherent of classical logic and bivalence, I believe that sorites sentences are false and that "sharp boundaries" claims are true. On any sorites series for a vague expression, I believe that somewhere in the series (not where we're looking) there is an object that possesses the property expressed by an utterance involving a vague expression right next to an object that lacks that property. I am reluctant, however, to call the proposed boundary between the property possessor and the property lacker a *sharp* boundary, since as I have stressed, this is but a metaphor, and I have as much right to the metaphor as does the proponent of gaps or degrees. I would cash out the metaphor in the following way: the boundary between the possessors and the lackers in a sorites series is not sharp in the sense that we can never bring it into focus; any attempt to bring it into focus causes it

to shift somewhere else. This essay has been devoted to developing and defending this idea, and to explaining why it provides the means for saying why we cannot find boundaries in a sorites series, and why we tend to believe of any given point in the series that the boundary is not there.

I have not, however, in any way addressed what I called the Semantic Question concerning the sorites paradox. That question was: if so-called "sharp boundaries" claims are true, how is that compatible with a vague predicate's having borderline cases? For "sharp-boundaries" claims seem to deny just that. If there are no truth-value gaps or degrees of truth between truth and falsity, then what is it for something to be a borderline case?

When trying to answer the question what is it for something to be a borderline case of a vague predicate, I think it important to begin by saying just what phenomenon it is we are trying to account for. Some philosophers, such as Fine, just define borderline cases as those things of which a predicate is neither true nor false. Timothy Williamson regards it as better to proceed by giving examples. The first approach will not do, since there are coherent accounts of vagueness incompatible with it. Williamson's approach will not do either, since it seems impossible to find examples about which people can agree.

We are prompted to regard a thing as a borderline case of a predicate when it elicits in us one of a variety of related verbal behaviors. When asked, for example, whether a particular man is nice, we may give what can be called a *hedging* response. Hedging responses include: "He's nice*ish*," "It depends on how you look at it," "I wouldn't say he's nice, I wouldn't say he's not nice," "It could go either way," "He's kind of in between," "It's not that clear-cut," and even "He's a borderline case." If it is demanded that a "yes" or "no" response is required, we may feel that neither answer would be quite correct, that there is "no fact of the matter."

In asking what it is for something to be a borderline case of a predicate, I think we should ask what might prompt one of this array of responses. There is no justification for assuming at the outset that it is always the same cause in every case. In fact, I think that hedging responses may have a variety of causes, which for the most part are to be counted as one or another form of ignorance.[29] The interest-relative theory of the meaning of vague expressions not only makes sense of this, but predicts it.

If asked to judge whether or not a certain car is *expensive* (for a car), what we must judge is whether or not that car costs a lot for a car, whether, that is, its cost is significantly more than it is typical for a car to cost. There is a lot of room for ignorance here. For one, we may have no idea of what the typical cost of a car is, or we may have merely inexact knowledge (in Williamson's sense) of what the typical cost of a car is. We may also have inexact knowledge of what our own interests are, of what exactly it would take for the difference between the car's cost and the typical cost to be a difference that is significant to us.

There are also more distinctly semantic (or if you prefer, pragmatic) sources of the ignorance which better explain why a hedging response might be provoked. It is convenient to place these under the heading 'ignorance of the context'. If the comparison class has been left implicit, for example, we may not know which comparison class is intended by our questioner. Even if it is clear that the comparison class is cars, rather than, e.g., essential possessions, we may still be unsure of which of the many norms of cost for a car is under consideration—are we to judge whether the car costs significantly more than a car *ought* to cost? Or are we to judge whether the car costs significantly more than the average car costs? Or are we to judge whether the car costs significantly more than anyone should be willing to pay for a car? Perhaps we are to judge whether the car costs significantly more than what cars used to cost when the price of cars was (as it seems to us now) more reasonable. The answers to these questions do not stand or fall together.

Yet another example is this: if we are joining in on a conversation already underway about whether the given car is expensive, there is room for even more ignorance, since in assessing whether the car costs significantly more than is typical for a car to cost, we may still feel unsure of what the right answer is to the question significant to whom? Are the interests of one but not the other of the conversants at issue? Is it both of their interests? Is it the interests of anyone who might want to buy the car? If the parties are in disagreement about whether the car is expensive, the reason we, as a third party to the conversation, might feel inclined to say, "There's no fact of the matter" (whatever that means) is closely related to the reason that some feel inclined to say that "there's no fact of the matter" about whether abortion is wrong. How the question is answered will depend on the interests or values of the person answering it, and we feel uncomfortable or unsure in giving the interests or values of one person special weight.

If ignorance, and especially "ignorance of the context," explains why we give hedging responses when we do, then a rejection of bivalence to account for borderline cases is unwarranted. Even if the idea that the truth-conditions of utterances containing vague expressions are both context-dependent and also sensitive to our interests in the way outlined here turns out not to be sustainable, my hope is that I have at least demonstrated that if we pay more careful attention to the way we actually use vague expression, there proves to be more room than is commonly thought within the space of classical logic and semantics to account for the many phenomena of vagueness.

NOTES

A kernel of the ideas in this essay first showed up in my "second-year paper" written at M.I.T. in 1994–95. Discussions with George Boolos during that year and the next had an enormous impact on my thinking on this and related topics. Diana Raffman's presentation of her work during her visit to M.I.T. that same year also had significant influence on my ideas about vagueness. I'm definitely in her debt. Discussions with Robert Stalnaker of the ideas that eventually showed up in another precursor of this essay—chapter three of my dissertation—really helped me to get clearer on exactly how context-dependence was (then) playing a role in my account of the sorites. Numerous discussions over the years about vagueness and context-dependence with Michael Fara, Michael Glanzberg, Richard Heck, Jason Stanley, and Timothy Williamson have been invaluable. Thanks finally to Christopher Hill and Timothy Williamson for helpful comments on the final draft.

1. As Kripke put it, whether or not a sentence leads to paradox does not depend "on anything intrinsic to the syntax and semantics" of the sentence. See Kripke (1975: 692).

2. Examples are Wright (1975: 329), Sainsbury (1991a: 6) and Soames (1999: chaps. 6, 7). Dominic Hyde (1994) defends the borderline-case conception of vagueness against the objections to it brought out here, on the grounds that 'is a borderline case' is itself a vague predicate. See Michael Tye's (1994) response for one worry about Hyde's account.

3. See also Sainsbury (1991b: esp. §VI).

4. I make an exception, however, for phenomenal versions of the sorites paradox—cases where 'F' is an observational predicate, e.g., 'looks red', and 'R' is an observational sameness relation, e.g., 'looks the same as'. With phenomenal versions of the sorites, I believe we should give up sentence (D). That is, if a clearly looks red and z clearly does not, I deny that a and z can be connected by a looks-the-same-as chain. See my "Phenomenal Continua and the Sorites" (Graff 2001) for a defense.

5. See §5 of Fine (1975) for his semantics for higher-order vagueness. The question at issue there is what is the correct semantics for the *definitely* operator we use to express vagueness—where to say that Herbert is neither definitely tall nor definitely not tall is to say that he is a borderline case. On the provisional semantics Fine proposes for the definitely operator, a sentence with an initial definitely operator (I'll use 'D' as an abbreviation) is always either true or false, with the result that "higher-order" predicates such as 'x is borderline tall' turn out to have no extension gap, and are therefore not vague. The more sophisticated semantic treatment of the definitely operator offered to accommodate the vagueness of higher-order predicates allows for sentences such as 'John is definitely tall' or 'John is borderline tall' also to lack a truth-value. What remains unchanged on the revised account, however, is that a sentence A will be true if and only if 'DA' is true. Thus even on the revised account, the extension of a predicate 'Fx' will be the same as that of 'D^nFx', for any number n of iterations of the definitely operator. Similarly, the extension of '$\neg Fx$' will be the same as that of '$D^n\neg Fx$' for any n. What I'm claiming we have no explanation for, on this account, is why we are unable to find the boundaries of these extensions. Sainsbury (1991b: §II) makes similar remarks.

6. See Fine (1975: 286).

7. See chaps. 7 and 8 of Williamson (1994).

8. One gets the impression that philosophers who are concerned about vagueness feel that those who accept bivalence are under some special obligation to provide an answer to the Epistemological Question. But I can't see why this should be so.

9. I'm assuming here that the domain is restricted to the members of some appropriately constructed sorites series.

10. See Williamson (1997), esp. §2.

11. Representative examples are Hans Kamp (1981), Manfred Pinkal (1984), Jamie Tappenden (1993), Diana Raffman (1994, 1996), Kees van Deemter (1996), Scott Soames (1999: chap. 7), and myself (1997).

12. See §III below for further discussion of Kamp's view.

13. See n. 11 above for references.

14. And even more so with the theory developed and defended in my dissertation (Graff 1997).

15. This must be distinguished from the frequently discussed point that adjectives such as 'tall', like 'good', are not intersective, in the sense that being in the extension of 'tall man' is not just a matter of being in the intersection of the extensions of 'tall' and 'man', since if it were, then from John's being both a tall man and a basketball player it would follow that he is a tall basketball player. In principle, adjectives could be non-intersective while still being extensional, and with few exceptions, are often treated this way. Wheeler (1972) and Cresswell (1976) serve as examples.

16. Well, it might be that as a matter of coincidence the things on my desk at a given time happen to form a kind. The point is that there is no kind which is such that something is a member of it at a time *in virtue of* its being on my desk at that time. When I say that we therefore have "no notion" of what a typical height is for a thing on my desk, I mean that there is *no* typical height for a thing on my desk, and we cannot make sense of the idea of there being one.

17. The fact that the variation in standards is subject to constraints is one indication that we should not merely liken such variations to Donnellan-style cases of misdescription. Alice Kyburg and Michael Morreau (2000) offer a nice discussion of this point.

18. I must express a special debt to Diana Raffman's (1994) work for inspiring examples like these.

19. Well, even that is not completely accurate. I am also presupposing that we have some relatively standard interpretation for the conditional, though not necessarily a material interpretation.

20. We should be wary of this result—that the collection of conditionals '$F(i) \to F(i+1)$' can be jointly true—given also that Kamp wants to maintain modus ponens as a valid rule of inference, and also that he does not adopt a nihilist position along the lines of that advocated by Peter Unger (1979). In order to avoid the apparent incompatibility of such a combination of views, Kamp proposes a nonstandard conception of validity. I refer the reader to Kamp (1981) for details. See esp. 260f.

21. Soames and I do give the same answer to the Epistemological Question, however. We both say that we cannot discover exactly which instance or instances of a given sorites premise are not true, because that would require locating the boundary (or boundaries) in a sorites series for the given vague predicate. But this is something we cannot do, since the boundary (or boundaries) can never be where we're looking.

22. See Soames (1999: 215).

23. See Raffman (1994: 47) and Raffman (1996: n. 18).

24. See Tappenden (1993: §5). I should add that Tappenden's account, though it does offer an answer to the Psychological Question, furnishes us with no answer to the Epistemological Question.

25. Some care would be required in further explicating what is meant by 'discriminate'. When I spill a few coffee grains, of course I *know* that the amount in the coffee scoop before the spill is not the same as the amount in it afterwards. My knowledge of the difference is not sufficient for me to count as discriminating between the two amounts, however, at least not in the sense intended here. If we take it that to discriminate between the amounts is to behave as if they were different, then we might lay down by stipulation that knowing is not behaving. Alternatively we might stipulate that the only behaviors under consideration are those directed at furthering my more or less immediate interests. I feel tempted here to echo remarks of Kripke's in *Naming and Necessity* by protesting that I am not offering a *theory*. It is not that I think that theories are bad in principle. I just want to acknowledge that some refinement or supplementation is definitely needed.

26. Heim and Kratzer (1998: §4.3) offer a nice discussion of this question with reference to some of the basic facts.

27. Kennedy actually proposes that the comparative morpheme '-er' is three-ways ambiguous depending on the form of the expression in the 'than'-phrase. See p. 133 (Kennedy 1999) for summary. The meaning I give here is the meaning associated with *phrasal* comparatives, as contrasted with comparative constructions in which there is more of a sentential-like component of the 'than'-clause.

28. It is proposed that the absolute morpheme, as well as other degree morphemes such as '-er' and 'too', occupy the head of an extended projection of an adjective. That adjectives have extended projections headed by degree morphemes is an independently motivated syntactic proposal due to Steven Abney (1987). The important point here is that although my proposal involves positing a rich semantic structure for superficially simple predications involving adjectives, correspondingly rich syntactic structures have been proposed on independent grounds.

29. I want to mention one source of hedging responses which does not stem in any way from ignorance and which also seems to have little to do with vagueness as such—namely, a desire to avoid unwanted conversational implicatures. What I have in mind is this: suppose you feel confident that a given stretch of road is not straight for a road; nevertheless, if it is not winding either, you may feel uncomfortable in asserting or unwilling to assert without qualification that the stretch of road is not straight (for a road), or if asked whether it is straight you may feel most inclined to hedge. In this case, it is not that you don't have a firm belief one way or the other about whether the stretch of road is straight. It is rather that in asserting without qualification that the stretch of road is not straight you would implicate that it is winding. We might account for the implicature by saying that typically your assertion that a certain stretch of road is not straight would have no point unless its departure from straightness were significant. I say that hedging with this sort of source has nothing to do with vagueness as such because it arises with precise expressions as well. You may know that *x* is taller than *y*, but still feel uncomfortable saying so without qualification, because in so doing you would (perhaps falsely) implicate that *x* is *significantly* taller than *y*, else your assertion would have little point.

REFERENCES

Abney, S. 1987. "The English Noun Phrase in its Sentential Aspect." Ph.D. thesis. Cambridge, Mass.: Massachusetts Institute of Technology.

Cresswell, M. 1976. "The Semantics of Degree." In *Montague Grammar*, ed. B. Partee. New York: Academic Press.

Fine, K. 1975. "Vagueness, Truth and Logic." *Synthese* 30: 265–300.

Frege, G. 1903. *Grundgesetze der Arithmetik, Begriffsschriftlich Abgeleitet*, vol. II. Jena: Pohle. Translation of selections taken from Geach and Black (1980).

Geach, P., and M. Black, eds. 1980. *Translations from the Philosophical Writings of Gottlob Frege*, 3d ed. Oxford: Blackwell. First edition published in 1952.

Graff, D. 1997. "The Phenomena of Vagueness." Ph.D. thesis. Cambridge, Mass.: Massachusetts Institute of Technology.

———. 2001. "Phenomenal Continua and the Sorites." *Mind* 110.

Heim, I. and A. Kratzer. 1998. *Semantics in Generative Grammar*. Oxford: Blackwell.

Hyde, D. 1994. "Why Higher-Order Vagueness Is a Pseudo-Problem." *Mind* 103: 35–41.

Kamp, H. 1975. "Two Theories about Adjectives." In *Formal Semantics of Natural Language*, ed. E. L. Keenan. Cambridge: Cambridge University Press.

———. (1981). "The Paradox of the Heap." In *Aspects of Philosophical Logic*, ed. U. Mönnich. Dordrecht: D.-Reidel.

Kennedy, C. 1999. *Projecting the Adjective: The Syntax and Semantics of Gradability and Comparison*. New York: Garland Press.

Klein, E. 1980. "A Semantics for Positive and Comparative Adjectives." *Linguistics and Philosophy* 4: 1–45.

Kripke, S. 1975. "Outline of a Theory of Truth." *Journal of Philosophy* 72: 690–716.

Kyburg, A., and M. Morreau. 2000. "Fitting Words: Vague Language in Context." *Linguistics and Philosophy* 23 (6): 577–97.

Montague, R. 1970. "English as a Formal Language." In *Linguaggi nella Società e nella Tecnica*, ed. B. Visentini et al. Milan: Edizioni di Comunità. Reprinted in Thomason (1974).

Montague, R. 1973. "The Proper Treatment of Quantification in Ordinary English." In *Approaches to Natural Language*, ed. J. Hintikka, J. Moravcsik, and P. Suppes. Dordrecht: Reidel. Reprinted in Thomason (1974).

Pinkal, M. 1984. "Consistency and Context Change: The Sorites Paradox." In *Varieties of Formal Semantics*. Proceedings of the Fourth Amsterdam Colloquium, Sept. 1982, ed. F. Landman and F. Veltman. Groningen-Amsterdam Studies in Semantics. Dordrecht: Foris Publications.

Raffman, D. 1994. "Vagueness without Paradox." *Philosophical Review* 103 (1): 41–74.

———. 1996. "Vagueness and Context-Relativity." *Philosophical Studies* 81: 175–92.

Sainsbury, R. M. 1991a. "Concepts without Boundaries." King's College London Dept. of Philosophy. Inaugural lecture delivered 6 Nov. 1990.

———. 1991b. "Is There Higher-Order Vagueness?" *Philosophical Quarterly* 41 (163): 167–82.

Soames, S. 1999. *Understanding Truth*. New York: Oxford University Press.

Tappenden, J. 1993. "The Liar and Sorites Paradoxes: Toward a Unified Treatment." *Journal of Philosophy* 90 (11): 551–77.

Thomason, R. H., ed. 1974. *Formal Philosophy: Selected Papers of Richard Montague*. New Haven, Conn.: Yale University Press.

Tye, M. 1994. "Why the Vague Need *Not* Be Higher-Order Vague." *Mind* 103: 43–45.

Unger, P. 1979. "There Are No Ordinary Things." *Synthese* 41: 117–54.

van Deemter, K. 1996. "The Sorites Fallacy and the Context-Dependence of Vague Predicates." In *Quantifiers, Deduction, and Context*, ed. M. Kanazawa, C. Piñon, and H. de Swart. Stanford, Calif.: CSLI Publications.

von Stechow, A. 1984. "Comparing Semantic Theories of Comparison." *Journal of Semantics* 3: 1–77.

Wheeler, S. C. 1972. "Attributives and Their Modifiers." *NOUS* 6: 310–34.

Williamson, T. 1994. *Vagueness*. London: Routledge.

———. 1997. "Imagination, Stipulation and Vagueness." In *Philosophical Issues*, vol. 8: *Truth*, ed. E. Villanueva. Atascadero, Calif.: Ridgeview.

Wright, C. 1975. "On the Coherence of Vague Predicates." *Synthese* 30: 325–65.