

The Epistemic Conception of Vagueness

Crispin Wright
University of St. Andrews

I.

According to the usual way of thinking, the indeterminacy on the margins of a vague concept is real. Vague concepts carve out their extensions rather as a blurred shadow carves out a region of the background on which it is cast. Thus, between the extension of such a concept and that of its complement, lies a blurry penumbral region—the domain of the borderline case. Vague concepts define the world rather as an imperfectly focused slide defines an image. Perfect precision, by contrast, is perfect focus.

This harmless-seeming imagery, however, seems prone rapidly to destabilize. Imagine that the blurred shadow and its background are quite large—apt for measurement in terms of tens of meters, say; and imagine a straight line run horizontally across them and calibrated in terms of centimeter points. Then it seems irresistible to say that no two adjacent such points are one in light and the other in shadow. But if shadow is conceived, as in the analogy, as if it were the conceptual complement of light, then that is to say that no point is such that, while it is in light, an immediate neighbor is not. And at least in classical logic, that is equivalent to the claim that if a point is in light, so is its immediate neighbor—which is all that is needed for a sorites paradox.

Perhaps the single largest program in contemporary work on vagueness consists in attempts, one way or another, to save something akin to the Fregean image of vagueness from paradox. One would think it must be possible—after all, blurred shadows and penumbral regions really do occur, so must admit, it would seem, of some form of distinctive yet coherent characterization. Any such approach may be viewed as *indeterminist*: it accepts the reality, however exactly it should be described, of the conceptual analogue of the penumbral region. According to proponents of the *epistemic* conception, by contrast, the Fregean imagery actually utterly mistakes the character of (what we take to be) vague concepts. There *is* no genuine indeterminacy, no region of borderline cases between the red and the non-red, the bald and the non-bald, the small

and the large. In truth, these distinctions are—and must be—completely sharp. Vagueness is rather a matter of *ignorance*. The various forms of reaction within a speech community that we are accustomed to take to be a reflection of indeterminacy—hesitancy when all relevant information is in, or irreducible conflicts, or perhaps a consensus that any clear verdict would be somehow improper—these various phenomena come about not because there is any genuine indeterminacy, or a lack of a “fact of the matter” but because speakers do not know any better, do not know where the sharp boundaries lie. Somewhere in the progression along any sorites series, however marginal the differences between its adjacent elements, will come a sharp cut-off point; a last case where the target concept applies, immediately succeeded by a case where it does not. It is merely that we do not—and cannot—know where this happens

This bizarre-seeming view has one clear attraction and one clear merit. The attraction is that it allows an acknowledgement of the linguistic phenomena which we associate with vagueness to sit perfectly comfortably alongside classical logic and semantics, and thus to provide the simplest possible dissolution of the sorites paradox. The merit is that of bringing out that the ordinary idea of genuine semantic indeterminacy is not itself a datum, but a *proto-theory* of data—the linguistic phenomena noted—which do not themselves constitute or clarify what vagueness, as conceived indeterministically, should be taken to be. The—at least initial—availability of the epistemic view brings it home that this is material for theoretical, philosophical description, and that much discussion has tended to bypass this essential task.¹

Even so, the epistemic proposal is apt to seem utterly unmotivated—would hardly seem better than a superstition were it not for the recent efforts of philosophers such as Timothy Williamson and Roy Sorensen² to take on at least some of the explanatory obligations which it must incur. Two such obligations are paramount. First, and most obvious, some sort of case must be made to make it credible that the sharp cut-offs that the epistemicist postulates really do exist. Second, and less obvious but equally pressing, an account is owing of *why* we cannot locate these cut-offs—why, indeed, we lack any clear conception of what it would be to locate them.

I think it fair to say that, given the technically cumbersome and variously unsatisfactory treatments of vagueness currently on offer from indeterminists, strong responses to each of these obligations would suffice to put epistemicism in a very competitive position. That would be a remarkable turn-about for the modern debate, which has been conducted almost exclusively within an indeterministic setting. In what follows I

shall accordingly review the best responses known to me on each count, all deriving from the work of Sorensen and Williamson. Then I shall outline what seems to me the most basic type of reason why the philosophical treatment of vagueness must dismiss these defenses of epistemicism, resourceful as they may be, as merely *tours de force*, and continue to struggle towards a coherent indeterminist conception of the matter.

II. FOR THE EXISTENCE OF SHARP CUT-OFFS: (1) WILLIAMSON ON BIVALENCE

Suppose it accepted that the Disquotational Scheme,

(DS) 'P' is true \leftrightarrow P

holds for every contentful indicative substitution for 'P'. And suppose the indeterminist characterizes (determinate) borderline cases for a given such sentence as failures of bivalence—as cases where neither that sentence nor its negation express truths. A demonstratively presented borderline case of red, for instance, will be a case where neither "That is red" nor "That is not red" are true. Since DS entails that the denial that 'P' is true is equivalent to the affirmation of 'Not P', and that the latter is equivalent to the predication of truth on 'Not P', we are immediately ensnared in contradiction. The epistemicist's preferred conclusion will be that, since the denial of bivalence is incoherent, so is the notion of a borderline case. Since vague concepts are characterized by the generation of borderline cases, it follows that no concepts are genuinely vague. QED.

This simple train of thought does a lot of work in Williamson's exposition in particular.³ If it is right, then there is no coherent option but to recognize that there *have* to be sharp cut-offs in sorites series.⁴

But of course a number of assumptions are in play, viz.

- (a) That the Disquotational Scheme is acceptable for vague substituends;
- (b) That any biconditional entails the biconditional which results from negating both its halves;
- (c) That borderline cases involve truth-value gaps.

Each of these could conceivably be denied. Which would it be best for the indeterminist to deny? My own view, to which I shall return shortly, is that (c)—the "gappy" conception of the borderline case—is a mistake in any case. But there is room for manoeuvre even if (c) is accepted.⁵

There are some fairly immediate moves to rehearse. To begin with, it may seem doubtful whether there could be any coherent unqualified rejection of the DS for vague sentences or any others. A sentence may be said to be true just in case the proposition it expresses is true. But the analogue of the DS for propositions, sometimes known as the Equivalence Scheme for truth,

$$(EQ) \quad \text{It is true that } P \leftrightarrow P$$

may seem uncontestable. If so, then a biconditional has to hold between any assertoric sentence and the sentence which says of that sentence that it is true. And that claim is just what the DS schematizes.

If that is accepted, the question becomes accordingly: what is the gist of that biconditional? Plainly, someone who accepts that a truth-apt sentence, *S*, may lack a truth-value cannot allow that *S* and the claim that it is true always enjoy the same truth-status—are always true, false, or “gappy” together. For when *S* is gappy, “*S* is true” ought, intuitively, to be false (since *S* determinately has some status other than truth). The most such a theorist can accept is only that whenever either is *true*, so is the other. Let it suffice for the validity of the *weak biconditional* of *A* with *B* that necessarily, whenever *A* is true, so is *B*, and vice versa, the validity of their *strong biconditional* demanding, in addition, coincidence in their truth-status in all cases other than truth. Then a believer in determinate truth-value gaps *can* accept the validity of the DS—but only as schematizing a range of weak biconditionals.

The upshot would be a qualified rejection of assumption (b). For weak biconditionals are not unqualifiedly contrapositive, if I may so put it. More accurately, suppose we distinguish corresponding weak and strong negations: a *strong negation* of *A* is true in the one case when *A* is false (it does not matter for present purposes what value it takes when *A* is neither true nor false). *A*’s *weak negation*, by contrast, is true when *A* has any truth-status other than truth. Then the principle:

$$\frac{A \leftrightarrow B}{\text{Not } A \leftrightarrow \text{Not } B}$$

holds good only if both ingredient connectives are weak or if the biconditional is strong. By contrast, if the biconditional is weak, but negation strong, then the validity of $A \leftrightarrow B$ will be consistent with, say, *A*’s being false in a case where *B* is gappy; when only the strong negation of *A* will be true, so that the lower biconditional, even weakly interpreted, will fail. Since, so we are supposing, our theorist accepts the DS only

as a weak biconditional, it follows that Williamson's purported reductio will have to be run in terms of weak negation. But when so run, it terminates in no inconsistency—to predicate truth of 'Not P' is not, when negation is weak, to say something inconsistent with 'P's being gappy.

Williamson's argument should accordingly be reckoned ineffectual unless it can be shown that an endorsement merely of the weakly biconditional DS is less than is demanded: that the reasons for endorsing the DS at all enjoin viewing its biconditional as strong. Is that so?

The biconditional would have to be regarded as strong if the sentences it connects—instances of 'P' and 'P is true'—were equivalent in content: expressive of the same proposition, or whatever. That would entail that their introduction into a single form of context would always preserve content, and it would hence be unintelligible how their negations could differ in truth-status. Manifestly, however, such sentences do not coincide in content. For one thing, they deploy quite different conceptual resources—it is possible fully to understand a sentence of the form, 'P is true' without understanding the mentioned 'P'. For another, their cross-substitution in certain types of context may generate changes in truth-conditions: compare, for instance,

If "red" meant green, "grass is red" would be true

and

If "red" meant green, grass would be red.⁶

This is not decisive, however. For Williamson's thought could just as well proceed not on the basis of the DS but via the Equivalence Scheme itself:

It is true that 'P' if and only if 'P'

which will generate contradiction in just the same way once harnessed to the supposition that a given proposition lacks a truth-value, so that both

It is not true that 'P'

and

It is not true that 'Not P'

are enjoined. And there are not the same obstacles as with the DS to the idea that instances of the two halves of the Equiva-

lence Scheme do express the same content; surely the thought that 'P' just is the thought that it is true that 'P'.

To this the defender of the truth-value gap conception of the borderline case must reply, I think, that those simply are not the same thought—precisely because when the proposition that 'P' is gappy, the thought that it is true that 'P' will not be a gappy thought but a false one. The "gappy indeterminist" must stick to it that even the Equivalence Scheme is valid only as a weak biconditional.

It remains to be disclosed if any instability lurks in that proposal. It is not clear that there is any.

III. DO BORDERLINE CASES INVOLVE TRUTH-VALUE GAPS?

The foregoing line of response to Williamson's argument is premised on an acceptance, with Williamson, of assumption (c). The response is of interest mainly because, if coherent, it shows how we might live with that assumption without falling rapidly into the incoherence Williamson expects. But actually there is nothing mandatory about (c)—indeed it arguably quite mistakes the kind of indeterminacy involved in the usual run of borderline cases. Reflect that we do not, in general, expect agreement among otherwise competent judges either about which the borderline cases of a vague predicate are, or in their verdicts on what others are agreed are borderline cases. On the contrary, the normal idea of a borderline case is one on which competent judges may unite in hesitation but about which they may also permissibly differ. If Jones is on the borderline of baldness, that will, of course, make it allowable if we each judge that he is borderline, but also allowable—at least in very many cases—if you regard him as bald and if I do not.

It is crucial to recognize that this phenomenon—of permissible disagreement at the margins—is of the very essence of vagueness, and that to leave it out of account is merely to miss the subject matter. I am not denying that there are often *definite* borderline cases of a given distinction. But wherever a stable consensus can be elicited that something is on the borderline between two concepts, that is merely an indication that we could, if we wished, employ a concept intermediate between them and hence that they are not really complementary. Even where such a consensus operates, however, the vagueness of the concepts concerned will surface in the permissibility of competent judges' coming to divergent verdicts on their margins. If this were not so—if divergence were never permissible—then some particular verdict would have to be mandated in each case, and the boundaries between the concepts concerned would have to be sharp.

This point accords ill with a conception of indeterminacy as failure of truth-value for each of two contradictory claims. For so to construe the borderline case is to conceive of it as having a status inconsistent with the truth of each, and hence is a commitment to regarding either polar verdict as *mistaken*. We therefore do better, I suggest, to try to conceive being borderline as a status consistent with both the polar verdicts: for an item to be a borderline case on the red-orange border is for it to have a status consistent both with being red and with being orange, (so not red), precisely because it is for that item to have a status under which it *has not been determined* whether it is red or not. If it has not been settled whether or not x is F , that cannot amount to x 's having a status inconsistent both with being F and with being not- F ; if it were, then matters would have been settled after all— x would be neither.

Vague predicates, non-epistemically viewed, are predicates associated with indeterminacy. But indeterminacy should be conceived as a matter of things having been left *open*—which requires consistency with each of the relevant polar verdicts in a particular case. It is because the truth-value gap conception, by contrast, requires borderline cases to occupy a status inconsistent with each of the polar verdicts that it is apparently open to the threat that Williamson tried to develop. If I am right, that is already a misconception, compounded by Williamson, of the way a non-epistemic conception of vagueness should proceed.

IV. FOR THE EXISTENCE OF SHARP CUT-OFFS:

(2) SORENSEN ON LIMITED SENSITIVITY

Sorensen's argument likewise purports to show that there is no stable option but to credit what we take to be vague concepts with sharp boundaries, and that any non-epistemic conception of vagueness is consequently incoherent.⁷ According to the epistemic conception, any sorites series for a vague predicate F will contain a determinate last F and a determinate first non- F . This will be true no matter how similar in relevant respects the adjacent elements of the series are. So vague predicates, as conceived by the epistemic conception, are, in a clear intuitive sense, of *unlimited sensitivity*: that is, there is no degree of change, however small, in those relevant respects which is always insufficient to change the status of an item in point of F -ness.

So far so good. It may now seem—and so Sorensen takes it—that one who rejects the epistemic conception has to be one who accepts that vague predicates are of correspondingly *limited sensitivity*: more precisely, that for each such predicate,

there will be some degree of change, u , in some relevant parameter(s) such that no possible pair of items, one a positive, the other a negative instance of the predicate, differ only to degree u or less.⁸

Sorensen's argument is now directly to the conclusion that this concept of limited sensitivity, and hence any concept of vagueness which incorporates it, is actually incoherent. The argument proceeds by a meta-sorites. Let F be any vague predicate, say "short," and u a degree of change in some relevant respect to which F is insensitive. As like as not, F will be sensitive nevertheless to changes of the order of, say, thousands of u . Sorensen accordingly constructs this illustrative argument:

- (1) A sorites argument concerning "short man" has a false induction step if the step's increment equals or exceeds ten thousand millimeters.
 - (2) If a sorites argument concerning "short man" has a false induction step if the step's increment is n millimeters, it also has a false induction step if the step's increment is $n-1$.
-
- (3) All sorites arguments concerning "short man" having induction steps with increments convertible to millimeters have false induction steps.⁹

The conclusion is tantamount to the unlimited sensitivity of "short." But Premise 1 is undeniable. So the opponent of epistemicism must either reject Premise 2 or dispute the validity of the argument. Sorensen is surely right that the latter course is no real option. But why would it be awkward for a believer in limited sensitivity to reject Premise 2?

Because, according to Sorensen, to suppose Premise 2 false is to suppose an exact threshold—though presumably an unknowable one—to the degree of limited sensitivity of "short." That, however, would be to go epistemicist at second order, and so to give the game away, since there is no conceivable principled motivation for the indeterminist's making such an exception.

There is some room for manoeuvre against this. Someone might want to query, for instance, Sorensen's implicit equation of rejection of Premise 2 with its (classical) denial. But his argument is easily strengthened so as to preempt such skirmishes, and to provide a nastier pay-off for the indeterminist than to force him into *ad hoc*-ery. Consider this principle:

- (4) If some sorites argument for F that works with a series each pair of adjacent elements of which differ by exactly

n u, contains a major premise to which there is a counterexample in the series, then some sorites argument for F that works with a series each adjacent elements of which differ by exactly n-1 u will contain a major premise to which there is a counterexample in its series.

This entails that there are no predicates of limited sensitivity, for it entails that there is no lower limit to the sensitivity of any predicate which has both positive and negative instances. But Premise 4 is entailed by the supposition that F is of limited sensitivity.¹⁰ So if F is of limited (but non-null) sensitivity, then it follows that there are no such predicates. So there are no such predicates. All predicates with both positive and negative instances are of unlimited sensitivity, and hence perfectly precise!

It merits remark that this argument depends on no distinctively classical principles. It is intuitionistically valid, and indeed will run in any logic which sustains conditional proof, reductio ad absurdum, and the usual quantifier rules.

It is another question what precisely it shows. One thing Sorensen would presumably want to claim for it is that—like the version he himself presents—it preempts the incoherentist response to the sorites offered by, among others, Dummett¹¹ and Unger¹², according to which vague predicates really are sorites-prone, sustaining appropriate major premises by virtue of genuine features of their sense.¹³ This response demands the limited sensitivity (tolerance) of vague expressions. Limited sensitivity is what it has just been proved no predicate has.

That the argument militates against incoherentism seems, however, a misconception: if anything, just the opposite seems to be true. What is shown by the argument is that no predicate can be of limited but *non-null* sensitivity. The problem is to combine the hypothesis of limited sensitivity with the claim that larger changes transform F's into non-F's and vice versa. However, the incoherentist should not want to make the latter claim: his point is precisely that there are no determinate F's and non-F's. If anyone is put in trouble by the argument, then it is the commonsensical indeterminist who essays to accept both the coherence of vague expressions—their possession of at least some determinate positive and negative instances—and their limited sensitivity.

That said, it must immediately be granted that an argument that reduced the options to epistemicism and incoherentism would be of the greatest interest. On reflection, however, this weaker contention, too, has certainly not been demonstrated. To credit the argument, in either version, with that significance depends upon the so far unjustified assump-

tion that to deny that a vague predicate is possessed of the sharp boundaries attributed to it by the epistemicist must be to attribute limited sensitivity (tolerance) to it instead: the assumption, indeed, that vagueness just *is* limited sensitivity. That identification must and will simply be denied by any indeterminist—be she a supervaluationist, degree-theorist, fuzzy logician, or whatever—who believes she has the resources to take principled exception to the major premises of sorites paradoxes.

V. DEFINITENESS

It is worth reflecting that drawing a principled distinction between vagueness and tolerance need not depend upon any heavyweight—e.g., supervaluational or degree-theoretic—semantic apparatus. The intuitive tradition of finding use for an operator of *definiteness* in the description of phenomena to do with vagueness easily makes the necessary initial distinctions. Williamson assumed a characterization of the borderline cases of *F* as items of which neither '*x* is *F*' nor '*x* is not *F*' hold good. Adjust that to: items of which neither '*x* is *definitely F*' nor '*x* is *definitely not F*' hold good, and read—in accordance with my earlier suggestion that borderline-case status is a matter of genuine indeterminacy, rather than occupancy of an estate inconsistent with both poles—the claim that *x* is not *definitely F* in such a way that it does not entail that *x* is not *F*. Then we can now express *F*'s lack of sharp boundaries in an appropriate sorites series as consisting in the truth of the claim:

$$- (\exists x)(\text{Def}[Fx] \ \& \ \text{Def}[-Fx'])$$

a principle which, unlike the negation of the epistemicist's ascription of a sharp boundary:

$$- (\exists x)(Fx \ \& \ -Fx')$$

is not at the service of a sorites paradox.¹⁴

The epistemicist can hardly deny that it is permissible to introduce such an operator. For an operator of knowability will, in his view, have both stipulated features: the borderline cases of *F* will indeed be cases such that they are neither knowably *F* nor knowably not *F*, and there will indeed be no entailment from '*x* is not knowably *F*' to '*x* is not *F*'. What he doubtless will, when challenged, deny is that such an operator has any coherent interpretation suited to the needs of indeterminism.

Sorensen, for his part, offers no explicit argument for that claim. Williamson, however, does express skepticism on the point. He writes:

... what more could it take for an utterance to be definitely true than just for it to be true? Given that it cannot be neither true nor false, how could it fail to be definitely true other than by failing to be true? Such questions are equally pressing with 'false' in place of 'true'. Again, 'TW is thin' is no doubt definitely true if and only if TW is definitely thin, but what is the difference between being thin and being definitely thin? Is it like the difference between being thin and being very thin? Can 'definitely' be explained in other terms, or are we supposed to grasp it as primitive?¹⁵

The first two questions come to: why does not the schema (DEF \rightarrow):

$$P \rightarrow \text{Def}[P]$$

hold generally good? It does not seem intelligible that there should be any way for an utterance to be true save by being definitely true—at any rate, there is no species of indefinite truth. But if that is decisive, then—since the converse principle is uncontroversial—statements and their "definitizations" will be generally equivalent, and the introduction of an operator of definiteness will be pointless.

This is quite a subtle matter and needs treatment in detail. But let me air some reflections on it and suggest a direction.

First, to scotch one natural friendly thought. It might be supposed that what the indeterminist requires is not to explain how (DEF \rightarrow) may fail by dint of some statement's being true while its definitization is not—maybe there can indeed be no explaining that—but to explain how such a pair can contribute differently to the truth-conditions of contexts containing them. That need not be the same thing, as we know from consideration of the weak biconditional earlier. But this is a blind alley. Remember that the situation under review is that of an indeterminist who proposes no truck with truth-value gaps or additional values. We therefore have just the classical truth-values in play. And in that case, it seems that if 'P' and 'Def[P]' are indeed weakly equivalent—if each is always true if the other is—they must extensionally coincide in all circumstances, there being only one other truth-status, viz. falsity, to occupy. So no extensional contexts differing only in that one embeds 'P' where the other has 'Def[P]' can vary in truth-status. Since negation is extensional, it appears that Williamson's point has to stick. Mustn't the indeterminist either concede that 'Not P' and 'Not Def[P]' always have the same truth-status—so that the definiteness operator is pointless—or resign himself to working with more than the classical truth-values—with gaps, or additional values?

Well, let's review the data of the problem. The indeterminist wants 'Not P' and 'Not Def[P]' to differ in their content in

such a way that (i) there is no entailment from the latter to the former—but it does nevertheless seem very intuitive that (ii) ‘P’ and ‘Def[P]’ should be, if not weakly equivalent, then at least equivalent in *some* sense. (iii) There are not to be truth-value gaps, or additional truth-values. (iv) The truth-conditions of negations are to be a function of those of the statements negated. How can these four constraints be met simultaneously?

We can learn from a simple case where they are. Consider any statement ‘P’, and the corresponding statement that ‘P’ is warranted. Such a pair will have the same acceptability conditions, yet there will in general be no entailment from the first to the second; and indeed their negations, consistently with the truth-functionality of negation and the currency of just the two classical truth-values, have different conditions of acceptability. The explanation is simple. Any pair of statements which necessarily coincide in cognitive status will provide an example of co-satisfaction of the four constraints provided their truth-conditions diverge. And the moral, I suggest, is that the indeterminist should look for a broadly cognitive interpretation of the definiteness operator. Such an interpretation may put us in position to affirm, what at first seems absurd, that ‘P’ and ‘Def[P]’ do indeed have different truth-conditions and that (DEF \rightarrow) may possess a true antecedent but false consequent.

Here is a first shot, intended just as an illustration of the territory in which the search for an account might begin. The governing idea is to construe “definitely” along the lines of operators like David Wiggins’ “There is nothing else to think but that”,¹⁶ or—more graphically—“No one in his right mind could doubt that.” In the typical run of cases, the clear cases of a vague predicate will be decidable—items which are definitely, or definitely not, red, or heaps, or bald will be effectively recognizable as such. Suppose such a clear, decidable case of a predicate F—“red”, say—and suppose someone who believes of it that it is not-F, not on the basis of testimony or inference, or groundlessly, but by an, as it transpires, abortive attempt at the relevant basic mode of decision (in this case, looking and seeing). There must be an explanation of the unfortunate upshot: maybe her vision is defective, or the light is bad, or her view is obscured; maybe she misunderstands the predicate. Whatever the correct explanation, it will advert to factors such that to know that they were in operation would be to have a reason to mistrust her opinion even without knowing what it was or knowing anything about the color of the object. Call an opinion *cognitively misbegotten* if a factor of this kind—a factor whose operation could be used to explain the subject’s falling into error, and to know of which would be to have a reason to mistrust the subject’s opinion—contributes to its genera-

tion. Then the suggestion is that, in the sense that concerns us, for 'P' to be *definitely true* is for *any appropriately generated opinion that 'Not P' to be cognitively misbegotten*.

Now, on this proposal, the thought that borderline cases may be marked by the permissibility of conflicting opinions will amount to the idea that associated with a typical vague predicate will be a range of cases for which neither of a pair of conflicting opinions need be cognitively misbegotten. About a borderline case of "red", you and I may hold respectively that it is, and that it is not red without any factor operating in the generation of either opinion such that someone who knew of it would have a reason to mistrust that opinion. Earlier I argued that the phenomenon of permissible disagreement at the margins is of the very essence of vagueness. In terms of the present proposal, we should refashion that suggestion: the basic phenomenon of vagueness is one of the possibility of faultlessly generated—cognitively un-misbegotten—conflict.¹⁷

Note that the epistemicist need not disagree with this, as far as it goes. For the epistemicist, the borderline cases of 'red' are all determinately but *inscrutably* red; naturally, therefore, should opinions conflict about such a case, no factors need be at work of the kind which would explain error, or conflict, in the *decidable*—definite—cases. Those are factors which would explain misfirings of the ordinary basic way of determining color. Since that way of determining color cannot engage inscrutable cases anyway, there is no a priori need to ascribe conflicting opinions about such cases of misfirings in its operation. Interestingly, then, on the proposed conception of definiteness, the epistemicist differs from the indeterminist not by rejecting the latter's conception of what a borderline case is but by *adding* something to it: the addition, namely, of the hypothesis of universal determinacy in truth-value: the principle of bivalence.

Call an opinion based neither on testimony nor inference, nor groundlessly held, a *primary* opinion. Then the suggested interpretation of 'Definitely P' amounts, to recapitulate, to this: any primary opinion that 'Not P' is cognitively misbegotten. How do matters stand with (DEF⁻), once the operator of definiteness is so interpreted? Evidently, the principle must fail if it is indeed a possibility for a pair of subjects to conflict in primary opinions concerning 'P' without either component being cognitively misbegotten. For in that case we have—just as intended—both 'Not definitely P' and 'Not definitely not P' holding true. If (DEF⁻) were good, we should therefore possess the materials, by a double contraposition, for a contradiction.

The situation should seem quite straightforward from the epistemicist's point of view. If the case is one on the borderline where 'P' nevertheless undetectably holds true, 'Def[P]' may nevertheless be false since a primary opinion that 'Not P' need

not, in the precise relevant sense, be cognitively misbegotten. The indeterminist, by contrast, cannot give quite that explanation of the failure of (DEF^{\rightarrow}) , since—unwilling to accept bivalence—he lacks the materials for an insistence that there are such undetectable truths. Nevertheless, it is his intended conception that there is real indeterminacy—that to be a borderline case is not to occupy a status inconsistent with either 'P' or its negation being true. So it has to be, if not required, then at least *consistent* to suppose either of the parties in an appropriately but faultlessly generated dispute over 'P' to be right—and hence, since both 'Not definitely P' and 'Not definitely not P' will hold in such a case, (DEF^{\rightarrow}) cannot be valid.

As both theorists, however, can readily acknowledge, it remains that, for any vague statement 'P', admitting of primary opinions in the sense outlined, any warrant for 'P' will be a warrant to suppose that a primary opinion that 'Not P' will be somehow cognitively misbegotten. So it can be agreed on all hands that there be no *identifying* a counterexample to (DEF^{\rightarrow}) ; in any circumstances in which 'P' may rightly be accepted, the acceptability of 'Def[P]' is assured. There is accordingly a clear form of equivalence—co-warrantability—between 'P' and 'Def[P]': all four constraints on the problem noted above are respected.¹⁸

These remarks are, to stress, only offered as preliminary pointers, albeit in an, as it seems to me, rather intuitive and obvious direction. Much would need to be clarified by an adequate treatment, including in particular the question of how this kind of proposal might be developed to accommodate higher-order vagueness—the vagueness of statements prefixed by the definiteness operator. Perhaps the most basic problem for the indeterminist is to characterize what vagueness consists in—to say what a borderline case is. It is also one of the least investigated. The epistemic conception should not be allowed to draw strength from this neglect. There is no cause to despair that the situation can be remedied.

VI.

FOR THE EXISTENCE OF SHARP CUT-OFFS: (3) SORENSEN'S CLONES

Sorensen has one other interesting and independent line of argument for the epistemic conception.¹⁹ I shall prescind from some of the detail of his presentation, considering what I take—I hope correctly—to be the nerve.

For all the vague properties which interest us we can envisage a process of gradual change whereby an item slowly acquires, or sheds such a property—gradually turns red, or goes

bald, or becomes tall. Imagine a pair of qualitatively absolutely similar items—Sorensen's clones—which undergo such a process: say the process of growing tall, having started short. Suppose they grow at exactly the same rate. But suppose Clone A starts growing before Clone B. Now, the following principle may seem quite compelling:

- (C) If an item undergoes some finite process of change, then had it started earlier and changed at just the same rate, it would have finished sooner.

An immediate corollary would be that if a pair of items undergo isomorphic processes from absolutely similar starting points, and if the processes proceed at exactly the same rates, then if one starts first, it finishes first. It follows that Clone A finishes the process of growing tall before Clone B. But to say that Clone A finishes first is to say that at some time, t , it has become tall—so is then tall—while Clone B has not yet become tall. Let H^A be A's height at t and H^B that of B. Then "tall" has a threshold in the interval between H^A and H^B . But this interval could be as small as you like—for we could have B start to grow as soon after A as you like. It follows that "tall" has a precise threshold. QED.

There is nothing the matter with this imaginative and resourceful argument except its basic premise, (C). Say that a process is *limited* by predicates, F and G , just in case its inception, and conclusion are, respectively, marked by the acquisition of F and G . Thus, the process of evaporation of water from a bowl, for instance, is limited by the predicates of the bowl: "contains less water than it did"—the process has started, *ceteris paribus*, just as soon as the bowl contains less water than it did—and "contains no water"—it is over as soon as the bowl contains no water. Likewise, the processes of the clones growing tall are limited by the predicates, "has grown taller" at their inception—A's process, for instance, has started just as soon as it is true to say of A that it is taller than it was—and "is tall" at their conclusion: each process is over just as soon as it is true of the clone that it is tall (though it may, of course, continue to grow). Then the point is merely that the principle (C) stands in need of qualification: that it is good only if the process in question is either limited by the application of a precise predicate at its conclusion, or—more generally—by a predicate *no more vague than* that which limits its inception.

What is it for a predicate to be more vague than another? The following limited case will do for our purposes. Say that F has a certain span of insensitivity, s , as measured in terms of some precise parameter, just in case no pair of items differing within that span and in no other respect can be such that

while one is definitely F, the other is definitely not F. Then G is more vague than F if each is sensitive to changes in the parameter in question but there is some s which is a span of insensitivity for G but not for F.

Consider principle (C) as applied to any process limited at its inception by a predicate F and at its cessation by a vaguer G. The counterfactual, "... had it started earlier ..." will invite consideration of a range of cases which will include some in which F first applies to the subject item at times at which it did not actually apply but which lie within the span of insensitivity, s, of G. But clearly of no such hypothetical process will it be true to say that it would have finished sooner than the actual process—since that would require that G apply to the item at a point when its condition would differ by less than s from its condition at various stages in the actual process at which it was definitely not G, contrary to the hypothesis that s is a span of insensitivity for G. It remains to reflect that the processes undergone by Sorensen's growing clones are limited by predicates—"has grown taller" and "is tall"—which are in exactly this case; indeed, "has grown taller than" is arguably a precise predicate.

Of course, this way with the argument needs the tolerance-free construal of the idea of imprecision, and the associated recourse to an operator of definiteness proposed in response to Sorensen's limited sensitivity argument above. It stands or falls with that proposal.

VII. WHY CAN'T WE KNOW WHERE THE SHARP CUT-OFFS LIE?

At this point, no unanswered argument to mandate the epistemicist's belief in universal sharp boundaries remains in play. In fact, though, Williamson for one takes the stance that epistemicism would be justified purely by the operational advantages of its conservatism:

Classical logic and semantics are vastly superior to the alternatives in simplicity, power, past success, and integration with theories in other domains. In these circumstances it would be sensible to adopt the epistemic view in order to retain classical logic and semantics even if it were subject to philosophical criticisms in which we could locate no fallacy.²⁰

and even presumably if the view was otherwise entirely unmotivated! Given this hard-nosed line, it comes as something of a surprise to find Williamson ready to acknowledge that epistemicism has *any* explanatory obligations. But acknowledge it he does:

For most vague terms, there is knowledge to be explained as well as ignorance. Although we cannot know whether the term applies in a borderline case, we can know whether it applies in many cases that are not borderline. The epistemic view may reasonably be expected to explain why the methods successfully used to acquire knowledge in the latter cases fail in the former.²¹

This is the second main explanatory obligation distinguished earlier: to explain why—granted the sharp boundaries in which epistemicism believes, and notwithstanding the shortcomings in the arguments so far reviewed for that belief—we seem barred from knowledge of where exactly they fall.

Williamson is alone among the epistemicists whose writings are known to me in taking this challenge very seriously—to his credit.²² His answer, very roughly, is that any broadly *reliabilist* conception of knowledge will require that knowledge always be cushioned by a *margin of error*. In order to know where the sharp boundary falls in a spectrum of patches ranging smoothly and barely discriminably from red to orange, I would have to know of some patch *j* that it is red while at the same time knowing that *j'*, its neighbor, is orange. But I cannot know that an item, *x*, is red unless my impression that *x* is red is a reliable indicator that it is so—which will be true only if *x* is flanked on *both* sides by patches that are red, and which therefore—since knowledge is factive—cannot be known to be orange. In short: I could not reliably practice, just on the basis of ordinary vision and memory, a perfectly sharp red/orange distinction. Hence, if there is such a distinction, I cannot know since I could not be reliable about precisely where it falls.

This is an elegant response to the problem. Is it sufficient?

Let us put on one side any reservation about the reliabilist conception of knowledge which drives Williamson's suggestion. The more immediate doubt is whether it in any case has the wherewithal to explain everything that needs explaining. Consider any sorites series for a predicate *F*, and let *j* be the last *F* element in the series, and *j'* the first non-*F* element. What Williamson has explained, if anything, is why we cannot know the conjunction, '*Fj* & *¬Fj'*'; that is, he has explained why we cannot know where the borderline *is*. But it also needs explaining why we cannot know where, in a large class of cases, the borderline *is not*. There are, on the epistemicist account, many conjunctions associated with the series in question whose truth-value we cannot know: each *k* lying close to but not immediately at the borderline will be associated, on the epistemicist view, with a determinately true conjunction, '*Fk* & *F'''*', or '*Fk* & *¬Fk'*', whose effect is that the borderline lies elsewhere but which we cannot know. What explanation of *that* is in the offing?

Clearly, the reasoning outlined for 'Fj & -Fj' has no direct application. That reasoning works by appealing to the idea of a margin of error to justify the principle that ' $K[Fx] \rightarrow Fx$ ', and to the factivity of knowledge to conclude therefrom that never ' $K[Fx \& -Fx]$ '. There is no way of reasoning similarly to the conclusion that truths of the form, ' $Fx \& Fx$ ', or ' $-Fx \& -Fx$ ', cannot be known for any x lying in the appropriate range.

What is true, of course, is that knowledge, *per impossibile*, of the true 'Fj & -Fj' would thereby place a subject in position to know all the problematical instances of ' $Fx \& Fx$ ', or ' $-Fx \& -Fx$ '. But Williamson has provided an explanation of the impossibility of knowledge of the latter only if knowledge of the true 'Fj & -Fj' would not merely suffice but is *necessary* for that knowledge—only if there can be no way to know any problematical instance of ' $Fx \& Fx$ ', or ' $-Fx \& -Fx$ ' save by knowing the true 'Fj & -Fj'. Without a vindication of that principle, his explanation is at best incomplete.

Lets put the point somewhat differently. Suppose it is in fact impossible to know any more than we do about the property which, on the epistemicist's view, is denoted by F . What Williamson's explanation requires, if it is to extend to explaining not merely the unknowability of the true 'Fj & -Fj' but that of each true problematical conjunction, is the truth of the following subjunctive conditional: *were* it possible to know more about that property—enough to know of certain of the problematical places where the cut-off does *not* come that it does *not* come in those places—then one would know enough to locate the cut-off. If this conditional fails, then it ought to be quite coherent to recognize for Williamson's reasons that there can be no knowing where the cut-off comes but to wonder whether some improvement in our information might not resolve *some* presently imponderable claims about where it does not come. And in that case, if the latter are indeed essentially imponderable, Williamson has not yet explained why.

There are other possible misgivings. It is a datum for explanation not merely that we cannot *know*—do not know what it would be like to know—where the cut-off comes, but that we have no conception of what it would be like to be able to *justify* any particular belief about where it comes. So Williamson's explanation cannot be fully satisfactory if it depends on features of the concept of knowledge which are not shared by justified belief and other relevant forms of epistemic relationship. But on the face of it the explanation does so depend: it depends on the factivity of knowledge.

Clearly no non-factive attitude, ' R ', can be subject to a margin of error in the sense Williamson defines for knowledge: it cannot be that whenever ' $R[P]$ ', then ' P ' holds in all sufficiently similar cases, since ' $R[P]$ ' will not ensure the truth of

'P' in the first place. So, as Williamson is of course well aware, something additional needs to be said. He proposes a variety of suggestions²³; but I should have thought that, just so long as 'R' is a notion with some significant internalist component, the not-very-sophisticated reflection should suffice that in the ordinary run of cases, our standards of justification find no significant distinction between the adjacent elements in sorites series. For instance, since the basic mode of color appraisal is to look and see, it is *ceteris paribus* absurd to suppose that we might justify incompatible color judgments about items that look exactly alike. That is a weaker principle than the tolerance of "may justifiably be regarded as red" with respect to indiscriminable difference²⁴; but any solution to the associated sorites paradox which disarms that tolerance principle should surely conserve the weaker one.

However, all this skirmishing may distract from what is most fundamentally unsatisfying about Williamson's proposal. Revert to the case of knowledge, and consider, say, a series of canes running from something below 5' 6" to something in excess of 6' 6" in length, each differing by 1/16" from those adjacent to it. Now if I am restricted to a merely visual assessment, any knowledge I can achieve of the lengths of the canes will certainly be subject to a margin of error in excess of 1/16". So we can explain à la Williamson why I will not be able to know the cut-off point for the predicate, "is less than 6' long" *by means of such assessment*. But, of course, it has not thereby been explained why—and indeed it is not true that—I cannot know the cut-off point at all. For in this case I can merely measure up. I know exactly what property is denoted by "is less than 6' long", and can recognize the appropriateness of other methods of assessment besides the visual, par excellence measurement, and apply them.²⁵ That is: the most that Williamson has done is to outline a form of explanation why there is no knowing where the cut-off for a vague predicate comes *provided* we are restricted to knowledge acquired by the means of assessment on which we actually rely in the application of that predicate. Nothing has been done to explain why knowledge or justified belief about the exact location of the cut-off is impossible *tout court*.

It might at this point be questioned whether that absolute impossibility has to be reckoned among the explananda. Are not the only clear data first, that we have no way of knowing where the cut-off lies if restricted to usual methods and second, that we have no inkling of other appropriate methods? But the question deserves to be pressed: Is there, in the epistemicist's view, or is there not an absolute impossibility here? If there is, it has not been explained by Williamson. Further explanation is still wanted of why it is impossible to determine the appropriateness of other methods of assessment

and bring them to bear. Such an explanation might proceed, of course, by appeal to the impossibility of our knowing precisely what property is denoted by a vague predicate.²⁶ But then an explanation will be owing of that. If, on the other hand, there is no absolute impossibility, then it has to be in principle possible to know more about the property denoted by a vague predicate sufficient to identify additional pertinent methods of assessment; and an explanation will now be owing of how this additional knowledge might be accomplished. Either way, Williamson's play with margins of error falls short of what, in recognizing that there is an explanatory obligation here at all, he implicitly undertook to provide.

VIII. EPISTEMICISM, REFERENCE, AND DELIBERATE APPROXIMATION

What kind of property in general does the epistemicist take vague expressions to stand for? Precise ones, of course. But what type of precise properties, and why? There is usually no difficulty in finding a precise comparative relation to order a sorites series. Such a comparative will be associated with a range of (relatively) precise properties, potentially at the service of the description of a sharp cut-off. Thus, "is taller than" can order a sorites series for "is tall/small", and is so associated with properties like being exactly 5' 11" tall. Is it the epistemic view that, e.g., "tall" denotes such a familiar precise property or not? If so, then—if we cannot know where the sharp boundaries lie—it must be unknowable which one; if not, then what would appear to be unknowable is *what* kind of property "tall" denotes: some precise property which cuts off sharply in any suitable sorites series, but is not a metric property of height.

Surely, though, the second option is a non-starter. "Tall" has to stand for a property with the features that anyone taller than someone who has the property likewise has it, and that anyone not taller than someone who lacks it, likewise lacks it. It is hard to believe that there are any precise properties which have these congruences on "is taller than" other than the familiar metric ones.

If that is right, then the epistemicist position in general must be that, while we can readily form a conception of the *kind* of property referred to by a vague expression, we do not know which such property is actually its reference. What is owing, therefore, is an account of what makes it the case that, e.g., "tall" refers—on a particular occasion of use, if you like: an assumption of context-dependence makes no difference to the essential point—to a particular metric property, *F*, rather than to any other as close to *F* as you please.

Epistemicists say nothing, or almost nothing about this. Understandably: what could there possibly be to say? The root and most fundamental criticism of epistemicism is that it has no account to offer of what constitutes *reference*. And this lack is not merely what we all lack—a philosophically watertight such account. Epistemicism cannot even *gesture* at an account of how the transcendent referential relations in which it believes might be constituted. What is it about our attitudes, or intentions, or practices with “tall”, or about the causation of our uses of the word, which might suffice to settle that—unknowably—it is exactly at 5’ 10 13/31”, say, that its lower threshold is reached? We are ineluctably ignorant, it seems, not merely of the alleged thresholds of vague expressions, but of any idea of the nature of the facts which constitute their reference to the relevant precise properties. I will come back to this.

Is there, in any case, any scope for epistemicism in every case where ordinary thought would impute some degree of indeterminacy? Counterexamples might be thought to be provided by the particles which we seemingly use to introduce vagueness deliberately: particles like: “approximately,” “roughly,” “about,” “almost,” “not quite,” and so on. Usually such particles are redundant. In most ordinary contexts—buying canes in the garden center, for instance—a more relaxed interpretation of the truth-conditions of what people say will prevail than is literally warranted. I say I want six 7-foot canes; but it will be good enough if the canes are *approximately* 7 feet in length. In most ordinary contexts, the approximate truth of what is strictly a precise statement will be good enough; and since it will be understood that this is so, there will be no need to avoid a precise formulation. The contexts where particles of approximation have work to do are ones in which a precise interpretation will otherwise prevail. For instance, in a context in which its removability through a 3’ 1”-wide door frame is about to be settled with tape measure, you had better express your *merely visual* assessment of the width of a sofa by saying, e.g., that it is *about* 3’ wide. In such a context the particle will have the effect of relaxing, to some unspecified extent, what will otherwise be taken as a relatively precise set of truth-conditions. In general, the role of such particles seems unquestionably to be to introduce some conveniently indeterminate degree of flexibility. But for the epistemicist, there can simply be no such thing as relaxing the truth-conditions of a precise statement in such a way—for there is no such thing as introducing a genuine indeterminacy. So what is the epistemic view of such particles—what do we use them to do?

The epistemicist has no option but to say that when a particle of approximation is added one precise statement sup-

planted by another. For instance, if I claim that Jones is roughly 6 feet tall, the epistemicist must so construe the truth-conditions of what I say that, for some fixed k and j , the coincide with those of the statement: Jones is more than 5 k and less than 6 j tall. Of course, I do not and cannot know what these values are. And it is utterly obscure, for all any epistemicist has said, what mechanism determines them. But that is not now my point. The question is rather how, on the epistemic view, the relaxation—which it is the whole point of the added particle to effect—can be seen as being effected.

The epistemicist will presumably have to elaborate his view along these lines. What a speaker tries to do by the relaxation is to allow that her statement can be accounted true in each of a range of cases when Jones' height differs from—is greater or smaller than—6. Consider every case in which a speaker would reckon her relaxed statement true. The greatest and smallest of these will be flanked by a range of cases about which she will likely hesitate whether or not her statement should still be regarded as true; and these in turn will be flanked by cases which she would regard as beyond what is allowed for even by the relaxed statement. So all the epistemicist has to do, it seems, is to ensure that the relevant values for k and j above effect cuts within the area of her hesitation. Provided the mechanism—whatever it is—which determines these values places them in this region—and provided speaker and audience tacitly know this—there will be something like the intuitively intended relaxation. It will be as if she had said:

Jones' height lies within a precisely bounded region of whose endpoints I can tell you only that they are respectively somewhere between (say) 5' 10" and 5' 11 1/2", and 6' 1/2" and 6' 2" respectively.

It must be essentially along these lines that the epistemicist should propose to handle such particles. For it is utterly obscure how motivation might be supplied for the concession that particles of approximation generate genuine semantic indeterminacy, which would not compromise the epistemic view across the much wider, ordinary range of cases, where an intuitive vagueness has nothing to do with the presence of such particles.

The question arises, however, what relation such an account of the content of a relaxed statement is supposed to bear to the semantic intentions of its speaker—the intention, that is, that the *approximate truth* of the precise statement be good enough. Obviously, it would be no good for the epistemicist to concede that the content of this intention *is* indeterminate—that it is, just as common sense suggests, associated with only

vaguely demarcated regions of fulfillment and non-fulfillment. For if that were granted, what objection could there be to the idea that the semantic role of "roughly" and its kin is simply to carry over the indeterminacy of such an intention into a linguistic content? The view has to be that, when I intend that you should stand roughly here, for instance, the demarcation of the range of cases in which you would comply from that in which you would not is *already perfectly precise*, even before it comes to speech. *I cannot but have a wholly precise intention*; the attempt to relax precision results only in the supplanting of an intention whose exact conditions of fulfillment I know by another of whose exact conditions of fulfillment, broader but no less precisely demarcated than those of the former, I can have no idea.

Maybe this upshot—that I have intentions and other attitudes whose exact conditions of fulfillment I do not and cannot know—may not seem too bizarre to those weaned on semantic externalism. I may intend to drink a glass of water, for instance, without knowing exactly what water is, or to avoid arthritis without knowing exactly what arthritis is. Sure, in both cases I can also knowingly give one perfectly adequate formulation of my intention; but then the epistemicist will presumably rejoin that my preferred formulation—that you should stand roughly here—is likewise a knowing, perfectly adequate formulation of a condition of fulfillment whose precise boundaries I do not know.

But there are evident differences. If the concept of water is correctly taken to be that of a natural kind, it is courtesy of the intentions of those who deploy it that it is so. We intend, that is, that there should be conceptual space for an account of what water essentially is—an account of whose details we may well be in ignorance—and we have a prior conception, however inchoate, of how such an account might be achieved. By contrast, no ordinary subject who knowingly possesses the sort of approximated intention we are concerned with will conceive of themselves as intending to introduce a corresponding conceptual space—space for an account, e.g., of where exactly "roughly here" begins and leaves off—nor pretend to any conception whatever of what it would be to even look for one. There is simply nothing to inform such a search—beyond the manifestly inadequate constraint that the boundaries not be so located as to infringe clear verdicts—and absolutely no intuitive preparation for the idea that there is anything to search for.

Ordinary forms of content externalism always proceed on the basis of an intentional abrogation of semantic responsibility. We, so to speak, knowingly delegate the determination of the content and boundaries of certain concepts to external factors. But the kind of determination postulated by the

epistemicist is nothing we bargain for in this kind of way, and would have to proceed by mechanisms of which we have no conception. Epistemicism simply helps itself to a rampant form of content externalism. It should acquire not one iota of plausibility from association with the moderate, local, carefully argued, externalist claims familiar in modern philosophy of language and mind.

Of course, the dialectical balance would be different if the arguments to enforce global sharp boundaries had worked. Then, we should be compelled to work for a conception of reference to redeem epistemicism's borrowing. But those arguments do not work, and there is every cause for skepticism whether any could. In this setting, the epistemicist proposal comes across not as a serious philosophy of vagueness but as an invitation to a kind of semantic mysticism.²⁷

NOTES

¹ Proponents of the enduringly fashionable supervaluational approach to vagueness, for example, typically simply help themselves to the idea that there are cases to which acceptable valuations may assign differing truth-values—no explanation is attempted of what distinguishes such cases from those—the definite cases—on which all acceptable valuations must agree.

² Roy Sorensen, *Blindspot*. Oxford: Oxford University Press, 1988: 199–253. Timothy Williamson, "Vagueness and Ignorance," in *Proceedings of the Aristotelian Society*. Supp. Vol. 66 (1992): 145–162; and *Vagueness*. London and New York: Routledge, 1994. The epistemic view also receives sympathetic treatment at the hands of J. Cargile, "The Sorites Paradox," *British Journal for the Philosophy of Science* 20 (1969): 193–202 (see also §36 of his *Paradoxes*. Cambridge: Cambridge University Press, 1979) and R. Campbell, "The Sorites Paradox," *Philosophical Studies* 26 (1974): 175–191. For further references, see Williamson, *Vagueness*. 300, note 1.

³ See Williamson, "Vagueness and Ignorance," 145–150, and *Vagueness*. 187 and following. Essentially the same line is run in §26 of Paul Horwich's *Truth* (Oxford: Blackwell, 1990): 80 and following; but Horwich seems to fight shy of an outright endorsement of epistemicism.

⁴ Note the (suspicious) strength of the argument: it implies that there cannot be genuine indeterminacy *even as a result of explicit intention*. What if we deliberately only partially determine the application conditions of a concept? Williamson (*Vagueness*, 213–214) envisages a word "dommal" whose use we determine by stipulating that it is to apply to all dogs and to no non-mammal. Suppose we intentionally refrain from saying anything else. Is a cat a dommal? The natural view, that the matter has been left indeterminate, is squeezed out by Williamson's argument—a consequence with which he is content. But suppose a range of candidates for a job, and a preliminary shortlisting meeting at which we decide that Jones will be interviewed and that no one without a Ph.D. will be interviewed, but then adjourn without any other decision. Can one agree with Williamson's reasoning that cats are not dommals without being willing to conclude that Dr. Smith will not be

shortlisted?

I will return to the matter of vagueness by intention and the problems, if any, which it raises for epistemicism in Section VIII below.

⁵ The discussion to follow may be compared with the concluding section of Peter Simons' response to Williamson, "Vagueness and Ignorance," 173–177.

⁶ This, of course, is to parse the consequents of these conditionals as, respectively:

... it would be the case that ["grass is red" is true]

and

it would be the case that [grass is red].

But that seems alright. Alternatively, consider any example of this shape. Suppose the actual F would not be the F, and would not be G, under C-conditions. Then the counterfactual,

Were it the case that C, the F would not be G

will have a reading under which it expresses a truth; but that is manifestly no guarantee that

Were it the case that C, "the F is not G" would be true.

⁷ See Sorensen, *Blindspots*, 246–252.

⁸ Limited sensitivity is thus to be equated with *tolerance* in the sense introduced in Crispin Wright, "On the Coherence of Vague Predicates," *Synthese* 30 (1975): 325–365.

⁹ Sorensen, *Blindspots*, 249.

¹⁰ *Proof*.

Suppose F is insensitive to differences of one *u*, and assume the antecedent of Premise 4, i.e., suppose a sorites argument for F that works with a series, S, each adjacent two elements of which differ by exactly *n* *u*, and let *k* and *k'* be an adjacent pair of elements in S which constitute a counterexample to the relevant major premise. And now consider any item *k** which differs by exactly *n*–1 *u* from *k* and by exactly one *u* from *k'* (and in no other respect from either). Construct a series, S*, whose *n*th element is *k*, whose next element is *k**, and whose every non-initial element differs from its immediate neighbors by exactly *n*–1 *u*. We have *Fk*. Suppose *Fk**. Then there are a pair of items, *k** and *k'*, differing but by a single *u*, such that *Fk** and not *Fk'*, contrary to the hypothesis that F is insensitive to such differences. So not *Fk**, and *k* and *k** must accordingly be a counterexample to the major premise for an S*-based sorites. Hence, there is a sorites argument for F that works with a series each adjacent elements of which differ by exactly *n*–1 *u*, and which contains a major premise to which there is a counterexample in its series. So the consequent of Premise 4 holds on the supposition of its antecedent; so Premise 4 is true.

¹¹ Michael Dummett, "Wang's Paradox," *Synthese* 30 (1975): 301–324.

¹² Peter Unger, "There are no Ordinary Things," *Synthese* 41 (1979): 117–154.

¹³ See Sorensen's remarks against Unger on p. 250 of *Blindspots*.

¹⁴ I am simplifying. There are issues here about the coherence of

higher-order vagueness—specifically, a case to be made that, if the Definiteness operator has certain seemingly not-unmotivated features, it may yet be possible to generate a sorites whose major premise is to the effect that F is higher-order vague. See Wright, “Is Higher Order Vagueness Coherent?” *Analysis* 52 (1992): 129–139, and, for criticism, Dorothy Edgington, “Wright and Sainsbury on Higher-Order Vagueness,” and Richard Heck, “A Note on the Logic of (Higher-Order) Vagueness,” in *Analysis* (1993): 193–200 and 201–208 respectively.

¹⁵ Williamson, *Vagueness*, 194–195.

¹⁶ David Wiggins, “Moral Cognitivism, Moral Relativism and Motivating Moral Beliefs,” *Proceedings of the Aristotelian Society* XCI (1990–1991): 61–86.

¹⁷ It is a further question whether such a conflict, once in the open, can be faultlessly *sustained*. Of course, it might be wrong publicly to *profess* either of the conflicting opinions. But would it be permissible to persist in *holding* it in the light of another’s apparently faultlessly-generated disagreement? Well, evidence of irreducible conflict is not always—not, for instance, in the moral sphere—received as a defeater of what one regards as admissible opinion. I know now that many opinions which I regard as unacceptable could be debated only to a stalemate by ordinary moral reasoning. But I shall not pursue the issue here. What the reader should note is that the proposal in the text only concerns faultless generation: one who construes the permissibility of conflicting opinions about borderline cases in this way has so far no commitment to the idea that such opinions *remain* permissible as conflict comes to light.

¹⁸ An additional constraint, of course, is that the converse of (DEF[→]) should hold—that ‘Def’ should be factive. But this should not give rise to difficulty. Here is one supportive line of thought, conditional on the hypothesis that the subject matter of ‘P’ and its kin is within primary reach, as it were; i.e., is open to competent, non-inferential, testimony-independent assessment—as are color, heaphood, baldness, and height. Suppose ‘Def[P]’. Then any primary opinion that ‘Not P’ is cognitively misbegotten. But the same cannot, consistently with the hypothesis, be true of ‘P’. So it has to be possible faultlessly to generate a primary opinion that ‘P’. Hence—since, by hypothesis, the same is not true of ‘Not P—P cannot be a borderline case. But nor can P be false—or, by the hypothesis, it would not be true that any primary opinion to that effect is cognitively misbegotten. So ‘P’ must be true. QED.

¹⁹ Roy Sorensen, “A Thousand Clones,” *Mind* 103 (1994): 47–54.

²⁰ Williamson, “Vagueness and Ignorance,” 162.

²¹ Williamson, *Vagueness*, 216.

²² The obligation is surely quite general: anyone who takes it that the truth-values of a certain range of statements are (potentially) unknowable ought to explain why. Platonism about number theory, for instance, and realism about the past have such obligations—albeit obligations which, it seems, they are easily able to meet. The platonist can build straightforwardly on the contrast between the finitude of our abilities and the potentially infinite range of epistemically independent consequences of statements involving quantification over all the natural numbers, and the realist about the past needs to point only to the contingency of the existence of effects distinctive of any particular event. Neither type of account would provide any justification for the idea that the relevant kinds of statement must be determinately true or false, of

course—that is a different battle the realist must fight another day. The question now is not to justify a belief in the relevant range of determinate, truth- or falsity-conferring states of affairs but to explain why, granted that belief, such states of affairs may conspire to elude detection.

²³ See Chapter 8, section 7 of his *Vagueness* and §6 of his “Inexact Knowledge,” *Mind* CI (1992), 217–242.

²⁴ The weak principle implies merely that if *x* and *y* look exactly alike, and if *x* may justifiably be regarded as red, then *y* may not justifiably be regarded as not red.

²⁵ This objection is quite compatible with the idea that measurement-based knowledge is also subject to a margin of error principle. I am assuming merely that the margin of error is smaller than 1/16”.

²⁶ We may, I suppose, discount the possibility of the property’s being identifiable but recognizably undecidable over the relevant range of cases.

²⁷ Thanks to Bob Hale, Mark Sainsbury, John Skorupski, Roy Sorensen, Charles Travis, and Tim Williamson.