# Vagueness Without Ignorance

Cian Dorr

August 18, 2003

Is a glass that is two-thirds full pretty full? We don't want to say 'Yes'; we don't want to say 'No'. This reluctance on our part seems very different in character and origin from our reluctance to answer 'Yes' or 'No' to questions like 'Is there intelligent life on other planets in our galaxy?'. A natural thing to say is that while in the latter case our reluctance is due to ignorance, in the former case it has nothing to do with ignorance: even someone who knew all the relevant facts wouldn't want to say 'Yes' or 'No' to the question 'Is a glass that is two-thirds full pretty full?' Characteristically, when $\underline{a}$ is a borderline case of the predicate '$\underline{F}$', we are motivated to avoid either asserting or denying the sentence '$\underline{a}$ is $\underline{F}$' by considerations that have nothing to do with ignorance. In the first two sections of this paper, I will try to show how this "no-ignorance theory" can be developed into an illuminating account of the nature of vagueness and semantic indeterminacy. The remainder of the paper will be spent addressing what I take to be the most important objection to the no-ignorance theory.

## 1   Why don't we say 'Yes' or 'No' to borderline questions?

$\underline{A}$ and $\underline{B}$ have instituted a simple signalling system. They explore the jungle independently, looking for fruit-bearing trees. When one of them finds such a tree, she makes a noise: either a hoot or a yelp. ($\underline{A}$ and $\underline{B}$'s vocal apparatus doesn't allow them to make any other sounds.) When one of them hears a hoot or a yelp, he comes to believe that the other has

1

found a fruit-bearing tree. Moreover, he takes the noise he hears as evidence relevant to the question how much fruit are on the tree in question: if the noise was a hoot, he favours hypotheses according to which the tree has more fruit; if it was a yelp, he favours hypotheses according to which the tree has less fruit. But even when this evidence is taken into account, his credences remain smoothly distributed: the hypothesis that the tree has $\underline{n}$ fruit never receives much more or less credence than the hypothesis that the tree has $n + 1$ fruit. We can represent this situation diagrammatically as follows:
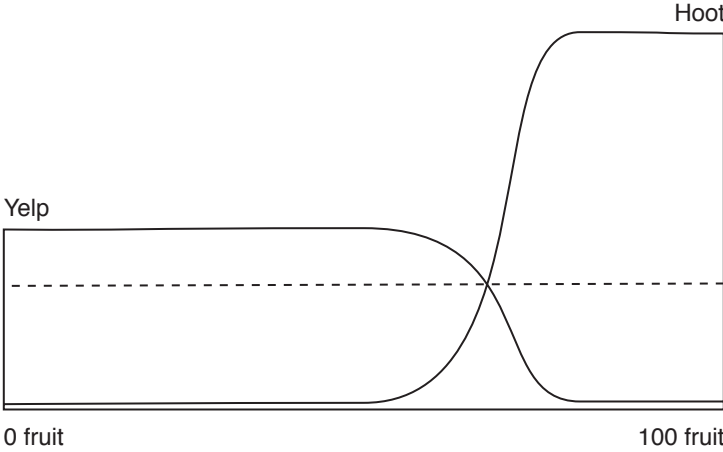
Here the dotted line represents $\underline{B}$'s initial credences (with regard to any given fruit-bearing tree that might be found by $\underline{A}$) and the two solid lines represent the credence-distributions he ends up with if he hears $\underline{A}$ yelp or hoot.

This state of affairs could persist even if we assume that $\underline{A}$ and $\underline{B}$ are perfect probabilistic reasoners and efficient decision-makers, wholly devoted to the goal of instilling in one another true beliefs about the distribution of fruit, and capable of determining with certainty exactly how many fruit are on a tree, and that each is perfectly confident that the other has all these features (and that the other is perfectly confident that $\underline{he}$ has all these features, and so $\underline{ad\ infinitum}$). Of course, if $\underline{B}$ knew all these facts about $\underline{A}$, and also knew exactly what probability $\underline{A}$ assigned to each hypothesis about how $\underline{B}$ would update his credences in response to a hoot or a yelp, $\underline{B}$ would not update his credences in the manner represented above; instead, his credences would end up looking like this:

Hoot

Yelp

0 fruit   100 fruit
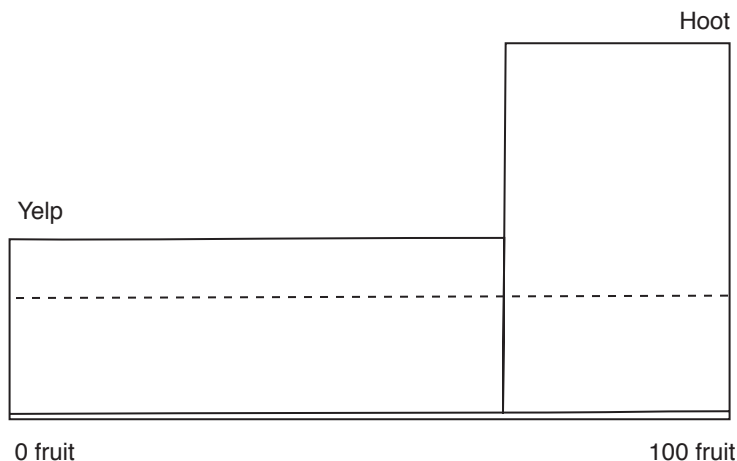
For the pattern of reactions depicted in the earlier graph to be rational, <u>B</u> must be uncertain exactly how <u>A</u>'s credences are distributed among the different hypotheses about <u>B</u>'s pattern of reactions.[1]

I boldly assert that the situation I have described is one in which <u>A</u> and <u>B</u> are speaking a very simple two-word language. Since the hoot and the yelp are clearly not precise words in this language, the language must be a vague one.[2]

As vague languages go, however, this one is quite exceptional, and not just because of its simplicity. For <u>A</u> has a decisive motivation either to hoot or to yelp whenever she finds a fruit tree: if she remains silent, <u>B</u> will not even know that she has discovered a fruit tree at all. In some cases this decision will be a very hard one, since both options carry a substantial and nearly equal risk of misleading <u>B</u>; nevertheless, the choice must eventually be made. And while the difficulty of the choice isn't due to any ignorance of <u>A</u>'s about the <u>tree</u>, the choice clearly would be much less difficult if <u>A</u> knew more about how each of the options would affect <u>B</u>'s credences.

But these features of <u>A</u> and <u>B</u>'s language-game are quite unstable: a small change in the situation would suffice to bring their practice much closer to actual vague language. Suppose that <u>B</u> acquires the ability to see for himself whether <u>A</u> has found a fruit tree, though he still cannot see how much fruit is on it. <u>A</u> will now have to consider a third option, that of remaining silent. Even if <u>A</u> thought that remaining silent would leave <u>B</u> with credences

3

distributed just as they would have been had <u>B</u> seen <u>A</u> finding a tree but been unable to hear her, <u>A</u> might prefer to remain silent if she was strongly motivated not to decrease <u>B</u>'s credence in the proposition stating the actual number of fruit on the tree in question. But since <u>B</u> can anticipate that the option of remaining silent will be tempting in this way to <u>A</u> in the "difficult cases", <u>B</u> will be inclined to take <u>A</u>'s remaining silent as evidence that <u>A</u> is actually faced with a difficult case, and hence to take <u>A</u>'s hooting or yelping as evidence that <u>A</u> is not faced with a difficult case. Since <u>A</u> can anticipate that <u>B</u> will react in these ways, <u>A</u> will now be strongly motivated to remain silent in such cases, since this will be the only way to avoid misleading <u>B</u>. So the result of our change to the environment will be to institute a new practice, in which <u>B</u>'s credences in case <u>A</u> hoots, yelps or remain silent will look something like this:
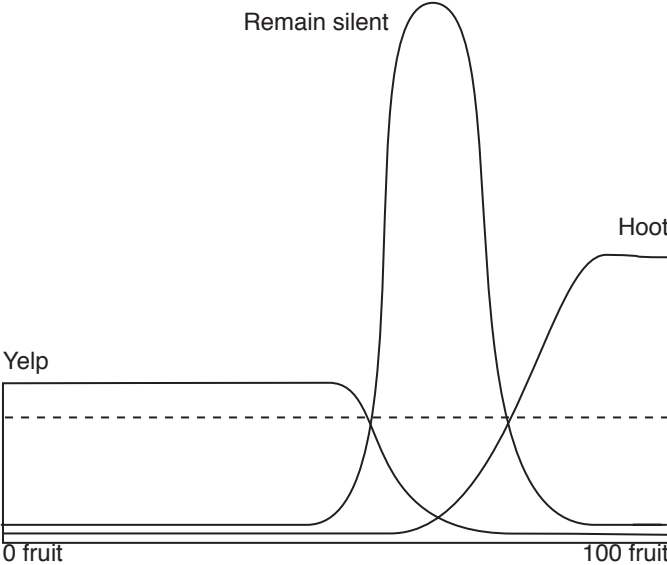

This new practice, unlike the old one, perfectly fits the no-ignorance theory. When <u>A</u> comes across a fruit tree that has, say, 65 fruit, <u>A</u> has a motivation to remain silent that has nothing at all to do with <u>A</u>'s ignorance of any facts.

    <u>A</u>'s choice whether to hoot, yelp or remain silent is analogous to the choice one faces when one has been asked a yes/no question: whether to say 'Yes', say 'No', or do something else (such as remaining silent). Consider the following situation:

Borderline Respondent can see a certain glass just well enough to know that it is between 60% and 70% full. Questioner, who cannot see the glass, asks 'Is the glass pretty full?'

Suppose, moreover, that both Questioner and Respondent are rational, competent speakers of English; that Respondent is a co-operative and honest person, strongly motivated not to mislead Questioner; and that these facts are common knowledge among Questioner and Respondent. Respondent can anticipate that if she says 'Yes' or 'No', Questioner will update his credences in something like the followinng manner:
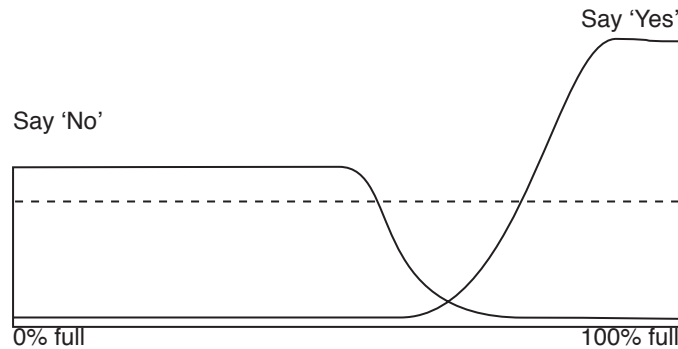
So in either case, Respondent expects that Questioner's credence in various important true propositions, such as the proposition that the glass is between 60% and 70% full, will be substantially lowered. This is just the sort of result that Respondent is most anxious to avoid. So Respondent will be strongly motivated to do something other than say 'Yes' or 'No'. She might, for example, keep silent. But in fact, this isn't an especially good choice, since it is apt to make Questioner think that Respondent has not heard the question, or that Respondent is not, in fact, a co-operative, honest, competent speaker. Thanks to the riches of the English language, Respondent has many better options: she could, for example, say 'It's hard to say', or 'It's around two-thirds full', or 'Sort of', or 'It's a borderline case', or 'I couldn't answer "Yes" or "No" to that question without misleading you'.

Borderline is paradigmatic of one sort of case in which we are strongly motivated not to answer 'Yes' or 'No' to a yes/no question. Here is a very different sort of case:

PRECISE   Respondent can see a certain glass just well enough to know that it
is between 60% and 70% full. Questioner, who cannot see the glass, asks 'Is the
glass at least 65% full?'

In this case the initial source of Respondent's motivation not to say 'Yes' or 'No' is the fact that these are <u>risky</u> options. If Respondent says 'Yes' and the glass is not in fact at least 65% full, or Respondent says 'No' and the glass is at least 65% full, Questioner will have been misled: his credence in various salient true propositions about the glass will have been lowered. Since Questioner will anticipate that Respondent will be motivated in this way, he will in fact be misled even if Respondent is lucky enough to guess right, since he will come to believe that either Respondent's vision is much more acute than it is, or that the level of water in the glass is further from 65% than it is. Because of this, the effects on Questioner's credences of saying 'Yes' and 'No' may be quite similar in PRECISE and in BORDERLINE; nevertheless, the origins of these effects are very different. One way in which this difference is manifest is in the appropriateness of Respondent's saying 'I don't know'. This is just the right thing to say in PRECISE; in BORDERLINE, by contrast, it would be quite a misleading thing to say, since it would be apt to make Questioner think, falsely, that Respondent can't see the glass very well.

In the remainder of this section I will consider two objections to this explanation of our unwillingness to answer 'Yes' or 'No' to borderline questions. First: this explanation only applies in situations where we really are primarily concerned to avoid misleading our interlocutors. But language, including question-answer exchanges, serves all sorts of other purposes beyond the communication of information. What explains our unwillingness to say 'Yes' or 'No' to borderline questions in situations where there is no chance that the questioner will be misled—in an exam, for example? My hope is that we can find explanations of our behaviour in such cases that are in one way or another parasitic on the explanation of our behaviour in cases where we do have to worry about misleading the questioner: in this sense I am committed to regarding language as <u>primarily</u> a tool for communication. In exams,

for example, we are typically motivated by the desire to convince the examiner that we are knowledgeable about the subject-matter and competent in the use of the relevant parts of the language. Saying 'Yes' or 'No' in a variant of BORDERLINE in which Questioner is an examiner who can see the glass as well as Respondent can would serve this goal poorly, because Respondent knows that Questioner expects her to behave <u>as if</u> she were really trying to convey information to Questioner, and will assess her knowledge and competence accordingly. If she says 'Yes', for example, he will conclude that she would have said 'Yes' even if she had been trying to convey information, and will infer from this either that she is bad at telling how much water is in the glass (if this is an eyesight exam) or that she is unfamiliar with the facts about the use of the expression 'pretty full' (if this is an English exam).

In other cases where we use language for non-communicative purposes, we have some motivation or other to say the first thing that comes into our head. This is true, for example, if we are philosophers or linguists trying to consult one anothers' "speaker's intuitions". To explain our failure to answer 'Yes' or 'No' in these cases, it suffices to explain why we don't have firm unreflective linguistic dispositions to give these answers. This explanation will advert to facts such as this: if we did have such dispositions, they would lead us to be systematically misleading in our communicative interactions with others; when we realised that this was happening, we would initially override our dispositions by reflection, and gradually lose the dispositions altogether.

Second objection: doesn't the explanation in BORDERLINE depend on Questioner's having a certain degree of mathematical sophistication? If Questioner were a young child, or a mathematical ignoramus, he wouldn't have the concepts necessary to entertain, let alone have degrees of belief in, propositions about the precise level of water in the glass.[3] Clearly this wouldn't make Respondent any more willing to say 'Yes' or 'No'. But what, in this case, would be the true propositions such that either of these answers would lower Questioner's degree of belief in them? We can hardly say that the propositions that the glass is pretty

full and its negation are both true, so that Questioner's coming to believe either of them to a lesser degree would constitute his being misled. We might try appealing to the proposition that it's indeterminate whether the glass is pretty full: but having to appeal to the concept of indeterminacy in this way in our explanations would undermine the project of giving a non-circular explanation of this concept in terms of facts about language-use.[4]

I can think of two ways to respond to this objection. First, I could maintain that mathematical sophistication is not, in fact, required for having degrees of belief in precise propositions about the level of water in the glass: such beliefs can truly be attributed even to animals. What mathematical sophistication brings is not the capacity to believe these propositions, but new and more explicitly articulated <u>modes of presentation</u> of the propositions. Alternatively, I could adopt a "Fregean" approach, and grant that the mathematically unsophisticated cannot literally speaking have degrees of belief in propositions like the proposition that the glass is at least 65% full: instead, they have degrees of belief in distinct but necessarily equivalent propositions, most of which aren't expressed by any English sentences. It doesn't much matter which of these options I take. What is important is that there is some sense in which we can consider various precise ways the glass might be and say things like 'It is more likely, from Questioner's point of view, that the glass is <u>this</u> way than that it is <u>that</u> way'. And it seems to me that we can make sense of claims like this, even applied to the most rudimentary thinkers. Moreover, it seems to me that thinkers—especially unsophisticated ones—generally have smoothly distributed credences, in the sense that when two possibilities are similar enough, one is only slightly more likely than the other, from the given thinker's point of view. Of course, these psychological claims are themselves subject to indeterminacy—the more so, the more unsophisticated the thinker—but that's not a problem for me. My explanatory ambitions would indeed be undermined if I had to use the concept of vagueness or indeterminacy in my explanation, but it's perfectly legitimate for me to use expressions which happen to be vague or indeterminate. Which is just as well, since otherwise I couldn't say much of anything.

2 Analysing semantic indeterminacy

It is not only in response to semantically indeterminate questions that we can find ourselves motivated not to answer 'Yes' or 'No' by considerations that have nothing to do with ignorance. A host of "pragmatic" factors can also give rise to the same sort of situation. Consider the following case:

> PRAGMATIC   Respondent can see that a certain glass contains just a few drops of water. Questioner, who cannot see the glass, asks 'Is there water in the glass?'

It would be misleading for Respondent simply to say 'Yes' in this situation. For Questioner probably thinks that Respondent probably thinks that the difference between a glass with just a few drops of water in it and one containing a substantial amount of water is of considerable relevance to Questioner's purposes; so Questioner will expect that if the glass does contain just a few drops, Respondent will say something more specific than a bare 'Yes'; so the effect of a bare 'Yes' will be to make Questioner conclude that there is probably a substantial amount of water in the glass. And of course it would also be misleading for Respondent to say 'No'. Nevertheless, this is not a case of semantic indeterminacy.

What is the relevant difference between PRAGMATIC and BORDERLINE? The following schematic answer—due in essence to David Lewis (1969, 1975)—seems to me to be promising. In PRAGMATIC, while a bare answer of 'Yes' would be misleading, it would not, unlike an answer of 'No', violate the conventions of language use—the conventions in virtue of which we count as speaking the English language.[5] In BORDERLINE, by contrast, there is no such asymmetry. Whether we should say that saying 'Yes' and saying 'No' would both be violations of the conventions, or that neither would be, at least we know enough about the concept of convention to see that both are on a par.[6] This is not to say that Respondent has an additional motivation not to say 'No' in PRAGMATIC in addition to her desire not to mislead Questioner, namely the desire to conform to the conventions of language. Rather, the fact that the conventions are what they are entails that our motivation not to mislead

one another is in most cases best served by conforming to the conventions.

If this is the right way to think about the difference between Borderline and Pragmatic, we should be able to analyse the notion of semantic indeterminacy in terms of the concept of convention. Here is one way in which such an analysis might be developed:

C1   A sentence S is underline{forbidden} for a speaker A at a time t in language L iff A at t has some property Q such that it is a linguistic convention constitutive of L that one should try not to utter S while one has Q.

C2 S is underline{permitted} for A at t in L iff A is not forbidden for A at t in L.

C3 S is underline{determinately true} for A at t in L iff S is permitted for A at t in L, and the negation of S is forbidden for A at t in L.[7]

C4 S is underline{semantically indeterminate} for A at t in L iff neither S nor the negation of S is determinately true for A at t in L.[8,9]

If these analyses are to have any hope of being successful, we will have to make some delicate distinctions as regards what does and does not count as a linguistic convention. Consider, for example, the sentence

(1)   There is no largest prime number.

One might suppose that the following rule has the status of a linguistic convention among English speakers:

(2)   One should try not to assert (1) unless one knows that there is no largest prime number.[10]

But if (2) counts as a convention, C1 entails that (1) is forbidden for any speaker who does not know that there is no largest prime number, and hence C3 entails that (1) is not determinately true for such a speaker. To avoid this absurd conclusion, we must deny that (2) has the status of a convention among English-speakers. The real convention—perhaps the only one—governing the assertion of (1) is as follows:

(3)   One should try not to assert (1) unless there is no largest prime number;

(2) is merely a strategy—a strategy any rational person would have good reason to adopt—for conforming to (3). (Or at least, if (2) does count as a convention, it is in some sense less basic than (3): so provided we restrict C1 so as to require that the conventions in question be "basic" in the relevant sense, it should be possible to avoid the absurdity.[11])

There is much more to be said about this analysis of semantic indeterminacy, and the notion of convention which figures in it. But I'm not going to press this inquiry any further in this paper.[12] I hope I have at least said enough to make you understand, if not share, my feeling that the no-ignorance theory contains the key to a proper understanding of indeterminacy and vagueness, phenomena whose place in reality would otherwise be utterly mysterious.[13]

3   Indeterminacy and knowledge

The objection to the no-ignorance theory which I will spend the rest of the paper discussing depends on the following premise:

> UNKNOWABILITY   Necessarily, if it is indeterminate whether $\underline{P}$, then no-one—at least, no ordinary human being—knows that $\underline{P}$.

Almost everyone who has written on the question seems to have found UNKNOWABILITY obviously true.[14] But UNKNOWABILITY, the claim that there is semantic indeterminacy of the sort attributable to vagueness (as opposed, say, to reference-failure), and the no-ignorance theory are jointly inconsistent. Anyone who accepts UNKNOWABILITY, accepts that vagueness sometimes leads to indeterminacy, and accepts the relatively uncontroversial logic needed to derive the contradiction, must give up the no-ignorance theory.[15]

To see how this works, let's look at what these claims entail about BORDERLINE. According to the no-ignorance theory, Respondent's unwillingness to answer 'Yes' in this case is not due to ignorance. I assume that if vagueness ever leads to semantic indeterminacy, it

is indeterminate in Borderline whether the glass is pretty full. So by Unknowability, we are committed to the claims that

(4a) Respondent doesn't know that the glass is pretty full, and

(4b) Respondent doesn't know that the glass is not pretty full.

Moreover, Unknowability entails that in any situation in which Respondent did know that the glass was pretty full, the glass would be determinately pretty full: in such a situation, she would have a much stronger reason to say 'Yes' than she actually does. So if, in the actual situation, Respondent is ignorant of the fact that the glass is pretty full, we can hardly say that her unwillingness to answer 'Yes' is not due to this ignorance. Hence (by modus tollens) it is not the case that Respondent is ignorant of the fact that the glass is pretty full. Similarly, since she would have a much stronger reason to say 'No' if she knew that the glass wasn't pretty full, it is not the case that Respondent is ignorant of the fact that the glass is not pretty full. But in general, the notions of ignorance and knowledge are connected as follows:

(5) $\underline{x}$ is ignorant of the fact that $\underline{P}$ iff $\underline{P}$ and $\underline{x}$ does not know that $\underline{P}$.[16]

So we have

(6a) It is not the case that: the glass is pretty full and Respondent doesn't know that
       it is pretty full.

(6b) It is not the case that: the glass is not pretty full and Respondent doesn't know
       that it is not pretty full.

But by some relatively uncontroversial logic (e.g. De Morgan's laws and disjunctive syllo-gism), (4a) and (6a) entail

(7a) The glass is not pretty full,

and (4b) and (6b) entail

(7b)   The glass is pretty full.

So if we accept UNKNOWABILITY, and accept that vagueness leads to semantic indeterminacy, we will have to give up the no-ignorance theory.

Perhaps that wouldn't be so bad if we could still hold on to the analysis of the notion of semantic indeterminacy which I presented in section 2. But in fact, there is no way to prise this analysis apart from the no-ignorance theory. For as I will now argue, the analyses C1–C4 are inconsistent with UNKNOWABILITY and the claim that vagueness sometimes leads to semantic indeterminacy.

The conception of semantic indeterminacy presented in section 2 would collapse if the concept of permissibility obeyed disquotational principles (with due allowance for speaker- and time-relativity), like

(8a)   The sentence 'That glass is pretty full' is permitted for A̲ at t̲ in English iff A̲ at
         t̲ is demonstrating a glass that is pretty full.

and

(8b)   The sentence 'That glass is not pretty full' is permitted for A̲ at t̲ in English iff
         A̲ at t̲ is demonstrating a glass that is not pretty full.

Given C3 and C4, principles like these would force us to give up the claim that there is any semantic indeterminacy in English, at least of the sort associated with vagueness. If there is indeterminacy of this sort, then presumably there could be a speaker A̲ and time t̲ such that

(9)   The sentence 'That glass is pretty full' is semantically indeterminate for A̲ at t̲,
         and A̲ at t̲ is demonstrating exactly one glass.

But, assuming that the sentences 'That glass is pretty full' and 'That glass is not pretty full' count as one anothers' negations, (8a), (8b) and C3 entail that

(10a) The sentence 'That glass is pretty full' is determinately true for $\underline{A}$ at $\underline{t}$ in English iff $\underline{A}$ at $\underline{t}$ is demonstrating a glass that is pretty full.

(10b) The sentence 'That glass is not pretty full' is determinately true for $\underline{A}$ at $\underline{t}$ in English iff $\underline{A}$ at $\underline{t}$ is demonstrating a glass that is not pretty full.

Substituting the right-hand sides of these biconditionals into C4, it follows that

(11) If the sentence 'That glass is pretty full' is semantically indeterminate for $\underline{A}$ at $\underline{t}$ in English, then $\underline{A}$ at $\underline{t}$ is not demonstrating a glass that is pretty full, and $\underline{A}$ at $\underline{t}$ is not demonstrating a glass that is not pretty full.

But this claim is clearly inconsistent with (9).

So if we want to accept C1–C4 and the claim that vagueness can lead to semantic indeterminacy, it is crucial that we reject disquotational principles for permissibility like (8a) and (8b). While there is nothing to stop us from granting that some of the conventions are "disquotational" in form, like

(12) One should try not to say 'That glass is pretty full' unless one is demonstrating a glass that is pretty full.

we cannot (given C1 and C2) say that all the linguistic conventions governing the assertion of sentences are like this. For example, we might say that one of the conventions governing this sentence is

(13) One should try not to say 'That glass is pretty full' if one is demonstrating a glass that is two-thirds full.

But if we accept UNKNOWABILITY, the classification of (13) as a <u>convention</u> will be unmotivated. The fact that we generally don't say 'That glass is pretty full' while demonstrating glasses that are two-thirds full can be accounted for perfectly well by (12), together with the non-linguistic fact that ordinary human beings can't know of glasses that are two-thirds full

that they are pretty full. Recall that the analysis of semantic indeterminacy depended on our having some way to explain why rules like this one should not count as conventions:

(2)   One should try not to assert 'There is no largest prime number' unless one knows that there is no largest prime number.

I can't see how a non-circular account of the notion of convention which succeeded in this task could fail likewise to entail that (13) is not a convention, assuming that UNKNOWABILITY is true.

Thus, as far as I can see, if we want to hang on to anything in the spirit of the no-ignorance theory, we have no choice but to reject (i.e. <u>refuse to assert</u>, not necessarily <u>assert the negation of</u>) UNKNOWABILITY. Of course, since '<u>x</u> knows that <u>P</u>' determinately entails '<u>P</u>', we must uphold the following weaker claim:

(14)   Necessarily, if it is indeterminate whether <u>P</u>, then no-one <u>determinately</u> knows that <u>P</u>.

But this is consistent with the claim that there can be cases where it's indeterminate whether <u>P</u>, and also indeterminate whether some ordinary person knows that <u>P</u>. This, I suggest, is what we should say about BORDERLINE, if we accept the no-ignorance theory: it is indeterminate whether Respondent knows that the glass is pretty full, and indeterminate whether she knows that it isn't pretty full.[17] And of course, similar claims are true of anyone who is in a predicament similar to Respondent's. Typically, when we know that it's not determinately the case that not-<u>P</u>, we don't determinately fail to know that <u>P</u>.[18]

Since knowledge implies belief, this entails that typically, when we know that it's not determinately the case that not-<u>P</u>, we don't determinately fail to <u>believe</u> that <u>P</u>. But it seems pretty clear that whatever it is in virtue of which we don't determinately fail to believe that <u>P</u> in these cases would be present even in cases where our belief that it's not determinately the case that not-<u>P</u> didn't amount to knowledge. So if we want to reject UNKNOWABILITY,

we will be naturally drawn to a picture on which typically, when we believe that it's not determinately not the case that P̲, we don't determinately fail to believe that P̲.

To my mind, this way of talking about belief seems quite natural. If the question in BORDERLINE had been 'Do you think the glass is pretty full?' rather than 'Is the glass pretty full?', I would be inclined to give the same sorts of "hedging" answers: 'Sort of'; 'In a sense'; 'I do and I don't'; 'Well, I think it's about two-thirds full'...[19] But I know well enough that many proponents of UNKNOWABILITY will have a very different reaction. To them, it will seem to be an "obvious datum" of introspection that when they believe that it's indeterminate whether P̲, they don't believe that P̲.[20] I'm not sure what I can say about this "argument from introspection", other than to encourage those who find it persuasive to think again about whether they might not be misdescribing what they find when they they look within themselves. What I can do, and what I will spend the rest of this paper doing, is to attempt to sway the balance of reasons by giving a new and independent argument against UNKNOWABILITY. Section 4 will be devoted to a defence of the central premise of this argument, the law of the excluded middle. The argument itself will be presented in section 5.

4   Excluded Middle

Suppose Questioner is trying to find out about the colour of a certain paint-chip from Respondent, who can see the chip; the chip is in fact borderline red-orange. Questioner asks 'Is it the case that the chip is either red or orange?' What should Respondent say in response, given her concern to avoid misleading Questioner? 'No' seems like a bad option. Assuming that Questioner's credences about the colour of the chip start out evenly distributed around the circle of hues, the likely result of Respondent's answering 'No' would be a credence-distribution looking something like this:
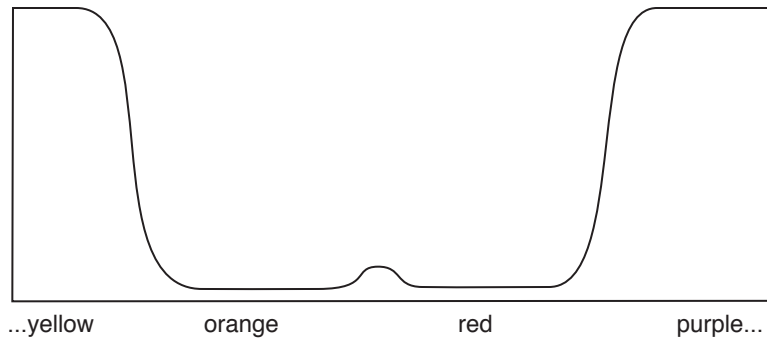
What about the option of saying 'Yes'? Here I find it especially hard to generalise with confidence about the reactions of typical English-speakers. I am inclined to think that saying 'Yes' is at the very least a <u>less</u> misleading option than saying 'No' in the circumstances. Nevertheless, it seems pretty clear that many speakers in many situations would end up with credences that look something like this:
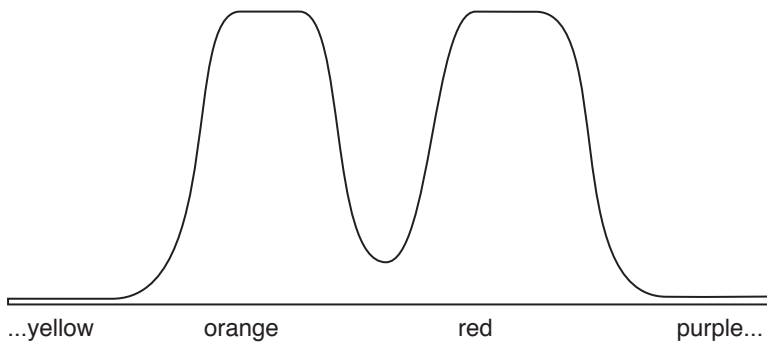
So Respondent has good reason to avoid giving either answer. Is this a genuine case of semantic indeterminacy, or is the difficulty of answering this question due to "merely pragmatic" factors? This is one of the most hotly contested issues in the philosophy of vagueness. Some say that the sentence

(15) The chip is either red or orange

is semantically indeterminate; others say it is determinately true but unassertable for pragmatic reasons. Those who say it's indeterminate will probably want to say the same thing

about all disjunctions with no determinately true disjuncts, as well as existential quantifications with no determinately true instances. So in particular, they will classify certain instances of the law of the excluded middle, like

(16) The chip is either red or not red,

as semantically indeterminate. Conversely, those give a pragmatic explanation of the unassertability of (15) will almost certainly want to explain the unassertability of (16) in the same way: for surely if (15) is determinately true, so is the claim that if the chip is orange, it is not red; and (16) seems to follow from (15) together with this claim. In this section, I will argue for a pragmatic explanation of the unassertability of sentences like (15), and by extension of instances of the law of the excluded middle like (16).

Consider the following arguments:

(17) (a) Every star with a surface temperature between 3000 and 5000 Kelvins is red or orange.

 (b) Arcturus is a star with a surface temperature between 3000 and 5000 Kelvins.

 (c) Therefore, Arcturus is red or orange.[21]

(18) (a) For every material object, there is a number that is that object's mass in kg.

 (b) Jupiter is a material object.

 (c) Therefore, there is a number that is Jupiter's mass in kg.

(19) (a) Whenever some people are all of different heights, one of them is shorter than any of the others.

 (b) The tall members of this department are all of different heights.

 (c) One of the tall members of this department is shorter than any of the others.

(17c), (18c) and (19c) seem to be unassertable in the same way as (15). On the other hand, there seem to be many contexts where it would be entirely appropriate, and not at all misleading, to assert (17a), (18a) and (19a). And (17b), (18b) and (19b) seem to be unproblematically assertable in almost any context.

How are we to resolve these paradoxes? I can see four options:

(i) The logically revisionist option. We could deny that the inferences are determinately truth-preserving in every context, claiming that, there are contexts in which, for example, (17a) and (17b) but not (17c) are determinately true.

(ii) The contextualist option. We could claim that sentences like (17a) and (17c) are in some way ambiguous: on one resolution, both are determinately true; on the other, both are indeterminate. The argument from (17a) and (17b) to (17c) is determinately valid provided the ambiguities are resolved uniformly. But typically, an assertion of (17a) would be interpreted as determinately true, while an assertion of (17c) would be interpreted as indeterminate.

An alternative version of this approach would hold that the sentences in question are context-sensitive rather than ambiguous. The argument is determinately valid if context is held fixed; but typically, the act of asserting (17c) would switch the context to one in which it is indeterminate.

(iii) The conservative option. (Conservative in its allocation of the status of determinate truth.) We could claim that sentences like (17a)–(19a) are semantically indeterminate, but assertable in many contexts for pragmatic reasons.

(iv) The liberal option. We could claim that sentences like (17c)–(19c) are determinately true, but unassertable in most contexts for pragmatic reasons.

Before we discuss the merits of these semantic claims, let's consider a psychological question: what explains the difference between our attitudes to (17a)–(19a) and our attitudes

to (17c)–(19c)? I tentatively propose the following hypothesis: we don't want to assert 'Arcturus is either red or orange' because this sentence naturally suggests the question 'Which is it, red or orange?', a question we know we will not be willing to answer straightforwardly by saying 'Yes', 'No' or 'I don't know'. Similarly awkward 'which?' questions are naturally suggested by (18c) and (19c).[22] By contrast, (17a)–(19a) don't have the same tendency to suggest "unanswerable questions". At least, this is so in those contexts in which these claims strike us as assertable: if the primary focus of our conversation has been Arcturus and its properties, (17a) <u>will</u> tend to suggest the question 'Which of these is the colour of Arcturus?', and will accordingly be unassertable.

It is an interesting question why we should be averse to asserting sentences which suggest unanswerable questions. It's not inevitable that this should be so: one can imagine a community for whom (15) was assertable even in borderline cases. Nevertheless, it's not just an arbitrary convention: it's the sort of practice that is apt to arise by a natural, rational process. The fact that a sentence naturally carries the mind to a question to which one would be unwilling to give a straightforward answer generates a fairly weak reason not to assert it: doing so will make it more likely that one's interlocutor will actually ask such a question, which is apt to sidetrack the conversation and lead to unpleasant practical dilemmas of the sort discussed in section 1. This reason in turn generates a second-order reason. Our interlocutor is likely, if he is smart, to anticipate that if the question suggested by a sentence is unanswerable in the context, we will be motivated by the first-order sort of reason not to assert the sentence. As a result, he will take our assertion of the sentence as evidence that this question would not, in this context, be unanswerable. Hence, if we assert the sentence, we will be providing our interlocutor with misleading evidence. And things will be even worse if our interlocutor anticipates the second-order reasoning as well as the first-order reasoning.... I'm not suggesting, of course, that we actually engage in all this complicated reasoning when we're deciding whether to assert something like (15). But seeing how rational considerations could foster the evolution of a practice like ours does, I think,

shed some light on the question why such practices should be so prevalent.

Here are three more pieces of confirming evidence for the "unanswerable question" hypothesis. First, consider two lists of sentences involving definite descriptions:

(20) (a) The mass of Jupiter is between 10^{27} and 10^{28} kg.

    (b) The mass of Jupiter in kg is between 10^{27} and 10^{28}.

    (c) The mass of Jupiter in kg is a number between 10^{27} and 10^{28}.

    (d) The mass of Jupiter in kg is one of the numbers between 10^{27} and 10^{28}.

    (e) One of the numbers between 10^{27} and 10^{28} is the mass of Jupiter in kg.

    (f) One, and only one, of the numbers between 10^{27} and 10^{28} is the mass of Jupiter in kg.

(21) (a) The first human being wasn't born until long after the dinosaurs were extinct.

    (b) The first human being was a creature which wasn't born until long after the dinosaurs were extinct.

    (c) The first human being was one of the creatures which weren't born until long after the dinosaurs were extinct.

    (d) One of the creatures which weren't born until long after the dinosaurs were extinct was the first human being.

    (e) One, and only one, of the creatures which weren't born until long after the dinosaurs were extinct was the first human being.

The sentences on each list appear to be logically or at least analytically equivalent. Nevertheless, while the first sentence on each list seems assertable, subsequent sentences seem more and more unassertable: it gets harder and harder to imagine a scenario in which one

could utter them without misleading one's audience. This is well accounted for by the hypothesis: each change of wording makes the question 'Which number <u>is</u> the mass of Jupiter in kg?' or 'Which creature <u>was</u> the first human being?' a bit more salient.

Second, observe that disjunctions with no determinately true disjuncts don't seem nearly so bad in cases where the speaker really is ignorant. Suppose, for example, that we are performing an autopsy on an alien being: finding that the aliens' retinas are sensitive to longer wavelengths of light than ours, I announce that

(22) The aliens' sun must be red or orange.

This remark of mine seems perfectly appropriate in the circumstances. Unlike 'The aliens' sun must be determinately red or determinately orange', my remark is unlikely to decrease my audience's credence in possibilities in which the alien's sun is borderline red-orange. This is explained by the hypothesis: in this context, the suggested question 'Which is it, red or orange?' <u>does</u> have a perfectly appropriate straightforward answer, namely 'I don't know'.

Third, many philosophers (e.g. Williamson 1994, pp.136–137) have noted that instances of the law of non-contradiction, like

(23) The chip isn't both red and not red

seem generally to be assertable. The same goes for sentences like

(24) It's not the case that the chip is neither red nor orange.

And yet, the only logical rules we need to derive the unassertable (15) and (16) from these sentences are De Morgan's Laws and (in the case of (23)) double negation elimination. This difference is neatly explained by the hypothesis: the forms of (23) and (24) don't naturally suggest any problematic 'Which is it?' questions. (Of course if you keep turning these sentences over in your mind, you can get yourself into a mood where unanswerable questions do come to seem pressing: if you do that, you'll find your willingness to assert the sentences falling away, as predicted by the hypothesis.)

Now that we have a grip on the psychological basis for our different attitudes to the premises and the conclusions of arguments like (17)–(19), let's return to the task of evaluating the four ways in which these facts might be accommodated in our semantic theory.

(i) I'll just say two things about the logically revisionary option, according to which the puzzling arguments have determinately true premises and indeterminate conclusions. First, this view rides roughshod over our intuitions about validity—intuitions which are not to be lightly dismissed, since it is largely by means of them that we uncover the compositional rules which explain our ability to understand and use indefinitely many sentences. Second, this view doesn't fit very well with the "unanswerable question" hypothesis. The extent to which a sentence suggests unanswerable questions is highly dependent on context: it is implausible that there is a single, context-independent point at which the sentences in (20) and (21), for example, switch from being determinately true to being semantically indeterminate. But we are going to have to introduce an element of context-sensitivity in any case, why not do so in such a way as to respect our intuitions about validity, as in the contextualist option?

(ii) The most important question faced by the proponent of the contextualist option is the question which words give rise to the relevant sort of ambiguity or context-sensitivity.

It might conceivably be suggested that the only sort of context-sensitivity we need is the familiar phenomenon of quantifier domain restriction. This lets us explain a few of the examples. For example, we could say that the contexts where

(18a) For every material object, there is a number that is that object's mass in kg

is determinately true are those in which the domain is restricted to include only precisely bounded objects, and exclude "fuzzy" objects which lack determinate mass. But this won't work in general, even if we don't mind the commitment to fuzzy objects. What sort of domain restriction, for example, could explain how

(17a) Any star with a surface temperature between 3000 and 5000 Kelvins is red or
orange

23

gets to be determinately true in certain contexts? It can't be that we are restricting the domain to stars to which our colour-words determinately apply, since that would allow (17a) to be true in this context even if some stars with a surface temperature between 3000 and 5000 Kelvins were borderline blue-green!

Another possibility is that the ambiguity or context-sensitivity of the problem sentences is due to ambiguity or context sensitivity in vague words like 'red'. I think I can see how this might work for sentences like 'the chip is either red or orange': we could say that 'red' and 'orange' have "penumbral connections" (see Fine 1975) on some disambiguations, or in some contexts, but not in others, and that phrasing things in such a way as to suggest unanswerable questions tends to make these words take on their disconnected senses. But how would this approach handle

(18c) There is a number that is Jupiter's mass in kg?

It is hard to see what sort of ambiguity or context-sensitivity in the words 'Jupiter' or 'mass' or 'number' could be responsible for this sentence being determinately true in some contexts and indeterminate in others.

The remaining possibility is that the ambiguity or context-sensitivity is due to logical particles like 'or'. In its simplest form, this approach would distinguish strong and weak senses of disjunction: strong disjunctions, unlike weak disjunctions, can be determinately true only when they have at least one determinately true disjunct.[23] Similarly, there will be a strong and a weak sense of 'some', of definite descriptions, and perhaps correspondingly weak and strong senses of 'and' and 'every'.

The main thing to be said against this is that words like 'or' just don't <u>feel</u> ambiguous or context-sensitive. We don't treat these words the way we treat ambiguous and context-sensitive expressions: for example, we have no hesitation in reporting someone who uttered the words 'the star is red or orange' as having <u>said that</u> the star was red or orange, without regard for changes of context. I don't think this sort of point is decisive: sometimes the right thing to do in response to a paradox is to recognise some hitherto unremarked sort of

24

ambiguity or context-sensitivity. But in general, we should strive to follow Grice's (1978) advice not to multiply senses beyond necessity.

If the logical particles really were ambiguous or context-sensitive, we would expect artificial languages in which this ambiguity or context-sensitivity is stipulated away by subscripting or some other such device to strike us as genuinely illuminating some sort of structure latent in natural language. But is this what we find? The logic of a language with distinct 'weak' and 'strong' versions of each of the logical particles is quite a strange world, with familiar classical rules and metarules being distributed between the different families of connectives in surprising ways. For example, given weak excluded middle

$$\models p \vee_{\text{weak}} \sim p$$

and the law of strong disjunction introduction

$$p \models p \vee_{\text{strong}} q$$

we cannot have the metarule of weak proof by cases

$$p_1 \models q$$
$$p_2 \models q$$
$$/p_1 \vee_{\text{weak}} p_2 \models q$$

since this would allow us to derive strong excluded middle from weak excluded middle. Investigating the properties of such hybrid logics certainly an interesting technical project, and may be of interest to reformers (like Field MSb) interested in improving on natural language. But it hardly seems plausible that anything like this is the true logic of natural language in its present state.

(iii) The big advantage of the conservative option is that it fits best with our intuitive response to the cases. If I assert (17a), and someone points out to me that it entails (17c), I will most likely react by backpedalling: I'll say something like 'I suppose all I <u>really</u> meant

was that every star with a surface temperature between 3000 and 5000 Kelvins has a colour somewhere in the spectrum between red and orange.'

The problem with the conservative option is that it is hard to see what sort of pragmatic mechanism could explain the assertability of the problem sentences. What rule or maxim could we be following when we figure out what the world would have to be like for an assertion of such a sentence not to be misleading, despite being literally semantically indeterminate? Whatever the rule is, it's going to have to somehow share or inherit the compositionality characteristic of semantic rules: we seem to able to compute the assertability-conditions of arbitrarily complex sentences of the problematic kind, and they seem to be logically very well-behaved, provided we avoid drawing consequences which too naturally suggest unanswerable questions. The challenge, once a pragmatic rule with this sort of complex character has been explained, is to say why it should be counted as belonging to the domain of pragmatics rather than semantics. Suppose that a proponent of this approach provides a translation-scheme that maps every English sentence $\underline{S}$ into another sentence S^*, and proposes that in certain contexts, a sentence $\underline{S}$ is assertable iff S^* is determinately true. If this proposal is to succeed in explaining the apparent logical relations among the problematic sentences, the operation $*$ is going to have to preserve logical structure. For example, it might involve replacing every constituent of the form '$\underline{P}$ or $\underline{Q}$' with 'Not determinately not ($\underline{P}$ or $\underline{Q}$)', or 'If not-$\underline{P}$ then $\underline{Q}$', and similarly for other connectives. But if anything like this is right—if, for example, there is a wide variety of contexts in which '$\underline{P}$ or $\underline{Q}$' is treated for communicative purposes just as if it were synonymous with 'If not-$\underline{P}$ then $\underline{Q}$'—why wouldn't it be better to say that in the contexts in question, '$\underline{P}$ or $\underline{Q}$' really $\underline{is}$ synonymous with 'If not-$\underline{P}$ then $\underline{Q}$?

This sort of problem becomes even more severe when we consider how the proponent of the conservative option might distinguish the contexts in which the assertability of a sentence depends on its literal, determinate truth from those in which assertability and determinate truth come apart. Given the evidence I have presented for the "unanswerable question" hypothesis, the natural thing to say is that the relevant feature of context is

whether a sentence naturally suggests an unanswerable question: if it does, it is evaluated literally; if it doesn't, its assertability depends on some other, compositionally specified condition. But this renders the question of literal semantic value idle: since sentences that do suggest unanswerable questions will always turn out to be semantically indeterminate, and hence unassertable, the only compositional rules we will ever actually need to use are those involved in the nonliteral mode of evaluation. We will never be called upon to use our literal understanding of the logical connectives in determining whether a sentence is determinately true. This seems quite bad: while we should of course recognise that assertability can come apart from determinate truth in many ways, we should at least expect the question whether a sentence is literally, determinately true to be relevant in many cases to speakers' deliberations about what to assert. Hence, the proponent of the conservative option is under considerable pressure to provide some other principle for distinguishing the two sorts of contexts: one on which, for example, the sentence 'the chip is red or orange' will be evaluated literally even if the speaker can clearly see that the chip in question is red. But the task of formulating a principle of this sort which agrees with the "unanswerable question" hypothesis in its predictions about assertability looks quite difficult. (How, for example, are we going to explain the assertability of (22)?) And even such a principle could be found, there would still be a concern that it wouldn't correspond to the actual psychological mechanisms underlying our deliberations in the way we should expect of a pragmatic theory.[24]

These challenges to the conservative option might be less worrisome if we had strong intuitions about the logic of 'or', 'some', etc., which were out of keeping with the allegedly nonliteral uses of these expressions in the problem sentences. Such intuitions could be regarded as reflecting our instinctive grasp of the most basic semantic rules governing these expressions, thereby motivating the treatment of the problem sentences as nonliteral. But in fact, the situation is the very opposite of this: it is the laws of classical logic that seem intuitive, when we consider them in an abstract light; it is only when we focus on particular examples like (16) that we are tempted to reject the classical laws. So, to the extent that

intuitions about validity are relevant to deciding where to draw the line between semantics and pragmatics, they tell in favour of the liberal view.

(iv) The main problem with the liberal option is the fact that it doesn't fit very well with our intuitive reactions to the cases. Sentences like

(15)   The chip is either red or orange

may start out by seeming assertable, but they come to seem more and more unassertable as we dwell on what they seem to be saying. According to the liberal view, this is a case where our first reactions are more reliable as a guide to literal, determinate truth than the judgments we arrive at after reflection. How bad is this problem? That depends on how we answer the much-disputed question to what extent the intuitive judgments of speakers can be relied on as a guide to semantic content (for a helpful overview, see King and Stanley forthcoming). There is an extremely optimistic answer according to which competent speakers can always tell, by consulting their speaker's intuitions, whether the unassertability of a sentence is due to semantic factors—i.e. to the failure of the sentence to be determinately true—or to pragmatic ones. At the other extreme, there are those who hold that the notion of literal semantic content is a construction of purely theoretical interest (if any), about which ordinary people don't even have opinions, let alone knowledge. In between, there is a continuum of moderate answers. The liberal option does indeed conflict with the extremely optimistic answer. But the required departure from optimism is not so great. It is <u>much</u> less severe, for example, than the departure envisaged by Millians who claim that the sentence 'The Babylonians believed that Hesperus was Phosphorus' is determinately true, but unassertable for pragmatic reasons. That sentence immediately strikes us as false; this reaction is quite stable, at least until we are exposed to some relatively complicated philosophical arguments. By contrast, our reaction to sentences like (15) is much more tentative and equivocal. I am inclined to think, then, that the conflict with intuition isn't such a bad problem for the liberal option.

Once we can see our way past this problem, there is everything else to be said in favour of this option. For the liberal view, the pragmatic rule which determines which sentences get to be assertable is easily stated: it is just 'Assert sentences that are determinately true, provided they don't too naturally suggest unanswerable questions'. This has the sort of simplicity we expect from a pragmatic rule: there's no mystery about why it doesn't belong in the semantics. Moreover, this rule fits naturally with the underlying psychology in a way that is not matched by the other options. It's not that we follow the rule in the sense that we first determine that (15) is determinately true, and then decide not to assert it because it suggests an unanswerable question. No: the sentence simply strikes us as something we wouldn't want to assert in the circumstances. But it is not implausible that the inaccessible mechanisms which generate these strikings have a structure which matches the rule: we have a disposition which, if left to its own devices, would lead us to find sentences like (15) assertable, but this disposition is <u>masked</u> or prevented from manifesting itself by our disposition not to assert anything which too naturally suggests an unanswerable question. By contrast, the more complicated methods of determining assertability suggested by the contextualist and conservative options seem unlikely to be reflected in any underlying psychological mechanisms. Finally, as I have already noted, the liberal view is consistent with our intuitions about validity, which I take to support classical logic: this is important, since intuitions about validity are arguably a more reliable guide to semantic content than intuitions about the truth of individual sentences.

I conclude, then, that the balance of reasons counts in favour of the liberal view. Sentences like (15) are determinately true, despite being unassertable; so are instances of the law of the excluded middle like (16).

In the next section, I will put this conclusion to work in arguing against UNKNOWABILITY. But the conclusion is also of independent interest, in that it suggests a way of fleshing out the picture of semantic indeterminacy presented in section 2. Once we have got used to accepting excluded middle even in the context of vagueness, it is but a short step to accepting that

all the theorems of classical logic are determinately true, and that all the inference-rules of classical logic preserve determinate truth. If we take this step, the way will be opened for an account of determinate truth in the style of supervaluationism. For if the rules of classical logic preserve determinate truth, there are some maximal classically consistent sets of sentences—call them the <u>admissible precisifications</u>—such that a sentence is determinately true iff it is a member of all of them. If we accept C3, the analysis of determinate truth in terms of convention which I presented in section 2, this means we can aspire to derive all the conventions governing specific sentences from one grand convention: assert only those sentences that are true (for you, at the present time) in each of $\mathcal{L}_1, \mathcal{L}_2, \ldots \mathcal{L}_n$, where $\mathcal{L}_1, \mathcal{L}_2, \ldots \mathcal{L}_n$ are "languages" in the sense of formal systems, which somehow or other assign a classical truth-value to each sentence (relative to a speaker and a time).[25] This, in essence, is Lewis' (1969) account of the content of the conventions of language.[26]

The one part of standard supervaluationist dogma that I think we should reject is the identification of truth with supertruth (truth on all admissible precisifications) and the commensurate identification of validity with preservation of supertruth. These doctrines cause all sorts of annoying trouble, and contribute nothing to the supervaluationist's explanations: they deserve to be scrapped.[27]

5   From Unknowability to epistemicism

We have concluded that instances of the law of the excluded middle are determinately true. Strictly and literally speaking, then, the glass is either pretty full or it isn't, although this isn't assertable in any ordinary context. With this conclusion in hand, let me return to the task of arguing against Unknowability. Can we really live with the view that sentences like these are strictly and literally true?

(25) This two-thirds-full glass is either pretty full or it isn't, but no human being can
        know which it is.

(26) One of the tall members of this department is shorter than all of the others, but we will never be able to find out who it is.

(27) Each person has a mass in grams (to the nearest gram), but we can never know what it is, no matter how accurately we weigh them.

To my ear, these claims seem obviously, deeply repugnant to common sense. But if we accept excluded middle and UNKNOWABILITY, we have no choice but to accept (25). And similarly for (26) and (27): the first conjuncts of these sentences seem if anything to be more assertable than the first conjunct of (25). Since we should accept excluded middle, and should accept the first conjuncts of (26) and (27), we must reject UNKNOWABILITY.

Some philosophers will not find this argument convincing. Epistemicists, like Sorensen (1988) and Williamson (1994), may have got so used to asserting sentences like (25)–(27) that they no longer see why the rest of us find them so bizarre. And there may be some philosophers (e.g. Schiffer 1999) for whom the counterintuitiveness of (25)–(27) is due entirely to the counterintuitiveness of their first conjuncts: if any of the members of this group were convinced by the argument for excluded middle in the previous section, they will no longer find (25)–(27) unacceptable. But the majority of philosophers have found epistemicism completely incredible, for reasons that can't be wholly due to epistemicists' acceptance of excluded middle and other theorems of classical logic, since other views which retain these theorems haven't met with the same reaction. What is it about epistemicism that prompts this incredulous response, if it isn't just the classical logic? Surely it's can't be the epistemicist's analysis of vagueness in epistemic terms: even if epistemicists claimed nothing more than extensional adequacy on behalf of their account, the view wouldn't look any less implausible. But once we set aside the analysis, what more is there to epistemicism than the serious and level-headed assertion of sentences like (25)–(27)?

Maybe it will be suggested that the problem with epistemicists isn't their propensity to assert things like (25)–(27) per se, but their failure to assert some additional thing which, if asserted, would somehow take the sting out of these claims. But what could this additional

thing be? Williamson suggests that the notion of ignorance is conceptually tied to the notions of truth and falsehood, so that one could avoid having to conclude from (25) that human beings are <u>ignorant</u> about the question whether the glass is pretty full if one went on to assert

(28) The sentence 'The glass is pretty full' is neither true nor false.

Williamson argues that (at least in this context) the notions of truth and falsehood are disquotational, so that (28) entails the contradictory sentence

(29) The glass neither is nor is not pretty full.

I find this argument against (28) completely convincing. But I also don't agree with Williamson that (28) would let us block the inference from (25) to the conclusion that we are ignorant whether the glass is pretty full, or indeed do anything else to mitigate the implausibility of (25). (28) is a claim about an English sentence. The claim that human beings are doomed to ignorance as regards whether the glass is pretty full, by contrast, has nothing at all to do with language: if it is true, it would have been true no matter what we had meant by the words 'the glass is pretty full'. So it's quite implausible that (25) needs to be supplemented by something like the denial of (28) before it entails anything about our ignorance.[28]

   For similar reasons, I don't see how it the counterintuitiveness of (25) would dissipate if we went on to assert that

(30) The sentence 'The glass is pretty full' is semantically indeterminate.

This claim entails that if someone <u>did</u> know whether the glass was pretty full, the task of communicating this knowledge to other English speakers would be harder than we might have expected. One natural way to do so, namely asserting 'The glass is pretty full', is ruled out by whatever it is that motivates us not to assert semantically indeterminate sentences. But these observations about language use do nothing, on their own, to render (25) less incredible.

Would (25) be easier to accept if we went on to assert

(31) It is indeterminate whether the glass is pretty full?

This raises a difficult question as regards the interpretation of object-language operators like 'it is indeterminate whether' and 'determinately'. My own view is that these operators are best understood as a sort of injection into the object language of metalinguistic predicates like 'is a semantically indeterminate sentence' and 'is a determinately true sentence'.[29] It's not quite clear what this means; but I think that it should at least have the consequence that (31) is no more effective than (30) at making (25) palatable. Other views about the meaning of the operators don't do any better, as far as I can see. There is, for example, the view that 'determinately' is simply primitive, in much the same way that some hold metaphysical necessity, or the Priorian past-tense operator, to be primitive. This view makes it mysterious why ignorance of indeterminate matters should be less objectionable than ignorance of determinate matters—why does this distinction matter any more than the distinction between ignorance of contingent matters and ignorance of necessary matters? There are also psychologistic views according to which 'it is indeterminate whether P' serves to express a certain sort of distinctive psychological state (or perhaps expresses the proposition that an idealised inquirer would be in such a state).[30] This seems more promising than the other approaches, but I still have doubts. How is being in the relevant state of mind supposed to take the sting out of sentences like (25)–(27)? When these sentences strike us as implausible, they also strike us as things that anyone who wasn't confused (or in the grip of a theory) would find implausible. So the proponent of this sort of approach owes us an argument that there is a state of mind that an unconfused person could be in such that (25)–(27) wouldn't seem objectionable to someone in that state of mind; and its hard to see how this argument could work unless there was some other way to make (25)–(27) seem palatable.

I conclude, then, that the implausibility of the epistemic view doesn't lie in what the epistemicists omit to say, but in what they do say, and in particular in claims like (25)–(27). Some opponents of epistemicism may have thought they could somehow get away with

making such claims provided they kept their fingers firmly crossed behind their backs; but they can't. When our theory entails that a sentence is strictly and literally, determinately true, we should be willing—having first made sure that we will not mislead anyone by unwanted pragmatic implicatures—to assert that sentence with the utmost seriousness. But we shouldn't be willing to assert (25)–(27) in this spirit. Since we should accept the law of the excluded middle, there is only one way to avoid having to do so: UNKNOWABILITY must be given up.[31]

Notes

[1] And therefore B̲ must also be uncertain exactly how A̲'s credences are distributed among the different hypotheses about the manner in which B̲'s credences are distributed among the different hypotheses about the way A̲'s credences are distributed among the different hypotheses about B̲'s pattern of reactions; and so ad infinitum.

[2] One moral we can draw straight away is that vagueness is not, contrary to a view that has been advocated by Stephen Schiffer (1998) and Hartry Field (2000, MSa), a psychological phenomenon requiring some sort of modification in the idealised picture of belief and action represented in standard decision theory. For all I have said, A̲ and B̲ could be angelic beings whose psychologies have exactly the structure represented in standard decision theory, and no more—perhaps they have little models of the space of possible worlds inside their heads, and they believe and desire things in virtue of the distribution of certain fluids over this space.

[3] This way of putting the objection presupposes that the proposition that the glass is pretty full is not itself a proposition about the precise level of water in the glass—for obviously mathematical sophistication is not required for believing the proposition that the glass is pretty full! This is controversial: many philosophers would hold that the proposition that the glass is pretty full i̲s̲ one of the propositions about the precise level of water in the glass, though it's indeterminate which one it is. But even those who hold this view might still maintain that the mathematically unsophisticated can only have degrees of belief in a few of the propositions about the precise level of water in the glass : the proposition that the glass is pretty full, and the proposition that it isn't; the proposition that the glass is full, and the proposition that it isn't, etc. (A view of this sort is suggested in Weatherson MSa.) This weaker claim would still be enough for the objection to go through.

[4] Besides, this explanation wouldn't work for hearers who don't have the concept of indeterminacy.

[5]Why doesn't the regularity that people generally don't just say 'Yes' to the question 'Is there water in the glass?' when the indicated glass contains just a few drops of water count as a convention of language use? For the same reason, I suppose, that the fact that people in this country generally drive <u>safely</u> on the right hand side of the road doesn't count as a convention. This fact is a consequence of a convention—that we drive on the right— together with a non-conventional regularity—that we generally drive safely. Likewise, the fact that we don't just say 'Yes' in situations like PRAGMATIC can be adequately explained as a consequence of the conventions of language together with certain non-conventional facts: that we generally expect each other to answer questions in a helpful way, that we are generally concerned not to mislead one another, that these facts are common knowledge, etc.

I'm not sure whether Lewis's theory of convention (1969) can give these results. If it can't, so much the worse for it.

[6]Which option should we pick? I'm tentatively inclined to prefer the latter (saying 'Yes' and saying 'No' are both permitted by the conventions), on the grounds that the answers 'Yes, it's about two-thirds full' and 'No, it's about two-thirds full' both seem more or less assertable despite the fact that 'Yes' and 'No' on their own would be misleading. Or at least, the answers 'Yes, it's about two-thirds full' and 'No, it's about two-thirds full' in BORDERLINE seem considerably <u>more</u> assertable than the answer 'No, there are a few drops of water in the glass' does in PRAGMATIC. The latter is liable to leave Questioner completely mystified, whereas we have little trouble taking the former in our stride.

This claim does not entail that the sentence 'the glass is pretty full and the glass is not pretty full' can permissibly be asserted in this or any context. Nor does it entail that one could permissibly assert 'the glass is pretty full' and then go on to assert 'the glass is not pretty full': the act of asserting the first sentence might change the context to one in which the assertion of the second sentence would be impermissible.

[7]You may object that an analysis of 'determinately true' should take the form of an analysis of 'determinately' together with an analysis of 'true'. I don't agree—I'm inclined

to think that 'determinately' should be understood as a sort of injection into the object language of a fundamentally metalinguistic notion (c.f. Fine 1975, p. 148). But you should feel free to read C3, like C1 and C2, as a stipulative definition made for the sake of the analysis of the metalinguistic predicate 'semantically indeterminate' in C4.

^{8}These analyses presuppose that S̲ has̲ a negation. What about a sentence in a language so simple that it doesn't have an operation of negation? I'm not sure what the best thing to say in such a case would be. Perhaps it's indeterminate whether such sentences are semantically indeterminate when they're not permitted?

^{9}C4 allows for two different ways for a sentence S̲ to be semantically indeterminate: S̲ and its negation could both be permitted, or they could both be forbidden. This seems like it might be a useful distinction. The first sort of indeterminacy seems to be found in the language of the fruit-counters from section 1, and arguably also in actual vague languages (see footnote 6 above). The second sort of indeterminacy might more plausibly be attributed to certain sentences involving expressions introduced by incomplete implicit definitions: for example, the word 'smidget', introduced by the stipulation that anyone under 4 feet is a smidget and anyone over 5 feet is not a smidget (Soames 1999).

^{10}One could run a similar objection substituting 'believes' for 'knows'.

^{11}A distinction between different kinds or levels of conventions might also be useful in making room for the phenomenon of conventional implicature (Grice 1975). One might think that 'Bush is a president but he is a politician' is determinately true, even though that it's a convention of language that one should try not to assert this sentence if there is nothing surprising in a president's being a politician. We could allow for this by saying that this convention is, in a sense that remains to be explained, less "basic" than the ones relevant to semantic indeterminacy.

^{12}One question I haven't talked about at all is the question how vagueness differs from other sources of semantic indeterminacy. Here is an idea that seems promising to me: the characteristic mark of a vague sentence is the fact that hearers will typically, when they

hear the sentence asserted, react by updating their credences along some dimension in the smoothly curving manner illustrated in the graphs in section 1, even when we factor out the effects of the hearer's doubts about the reliability and trustworthiness of the speaker.

I think it's plausible that vagueness, on this conception, always brings with it the epistemic possibility of semantic indeterminacy, as analysed in C4, and indeed higher-order indeterminacy. But I won't argue for these claims here.

[13]One feature of this account which people have tended to find surprising is its identification of semantic indetermincy as a distinctively public-language phenomenon. Couldn't there be vagueness in mental representation as well as linguistic representation? Well, of course I can agree that it could be vague whether someone believes that P, whether it visually appears to someone that Q, etc.—indeed, I will shortly be arguing that this sort of situation is much more pervasive than has often been thought—but those who press this worry seem to be looking for something more than that. What they want, I think, is an analysis of what it would be for a word of sentence in the language of thought to be vague, that doesn't merely refer this vagueness back to the vagueness of the (public) language we use to talk about the language of thought.

Perhaps I can oblige them without departing too far from the picture of vagueness I have been presenting. A sentence in my language of thought could have this feature: whenever it goes into my belief-box, it constitutively makes it be the case (other things being equal) that I believe that there is a glass in front of me, and that my credences over the various hypotheses about the amount of water in the glass are distributed in the same smooth way that Questioner's credences would be in BORDERLINE, if Respondent said 'Yes'. It doesn't seem like too much of a stretch to say that a mental sentence that played this sort of functional role would be vague, and indeterminate in those cases where the presence of either it or its negation in my belief box would decrease my credence in important true propositions.

[14]One exception is David Barnett (2000): while he does not explicitly reject UNKNOWA-

BILITY, he does make the similar claim that its being indeterminate whether P̲ does not rule out the possibility that even someone in ideal epistemic conditions with respect to the question whether P̲ could have good reason to believe that P̲. Another is Brian Weatherson (MSb), who presents an argument against UNKNOWABILITY based on the substitutability of certain kinds of expressions in knowledge-attribution contexts. (A similar argument is presented in Hawthorne forthcoming, though Hawthorne takes it to be a problem for the substitutivity principles rather than for UNKNOWABILITY.)

[15]This is not to say that they must accept the negation of the no-ignorance theory. They would be committed to this only if they accepted the rule of reductio ad absurdum, which is controversial in the context of vagueness, since it allows us to derive the law of the excluded middle. But having to give up the no-ignorance theory would be bad enough.

[16]Could we resist the argument at this point, by adopting an 'inflationary' conception of facts, on which 'it is a fact that P̲' is not analytically equivalent to 'P̲'? I suppose we could: but at the cost of rendering the no-ignorance theory trivial, by in effect helping ourselves to the crucial notion of determinacy which we were trying to explain. (C.f. Field MSa, pp. 1–3.)

[17]Since knowledge is factive, it's determinately the case that if Respondent knows that the glass is pretty full, the glass is pretty full, and if Respondent knows that the glass isn't pretty full, the glass isn't pretty full. But we don't have to endorse the converses of these conditionals: we are free to maintain that it's indeterminate whether Respondent is ignorant of the fact that the glass is pretty full, for example. To uphold the no-ignorance theory, we merely have to assert that determinately, if Respondent is ignorant of the fact that the glass is pretty full, her unwillingness to answer 'Yes' is not "due" to this ignorance: she has a reason for not answering 'Yes' that would be present even if she knew that the glass was pretty full.

It's just as well that we don't have to endorse the converse conditionals, since there's a good argument that they are false. Suppose the glass is in fact determinately 65.1% full. Determinately, Respondent doesn't know any proposition that entails that the glass is at

least 65% full. (She doesn't <u>believe</u> any such proposition; and even if she did, her belief would be too <u>unsafe</u> to count as knowledge.) But it's not determinately not the case that the proposition that the glass is pretty full is a true proposition that entails that the glass is at least 65% full. Hence it's not determinately not the case that the proposition that the glass is pretty full is a true proposition that Respondent doesn't know.

[18]Though not in all cases. Consider, for example, the name 'Princeton', which is indeterminate in reference between "Lesser Princeton" (Princeton Borough) and "Greater Princeton" (Princeton Borough plus Princeton Township) (this example is due to Lewis 1993). Suppose I know that exactly one of these entities is a municipality, but I'm not sure which it is. Then I know enough to know that it's indeterminate whether Princeton is a municipality, but we had better say that it's determinately the case that I don't know that Princeton is a municipality, since it seems to be determinate that none of the entities I have mentioned is such that I know that <u>it</u> is a municipality.

[19]Perhaps this fact about my reactions has something to do with my fondness for a view of belief which emphasises the analogy between the way belief represents and the way pictures represent (see Lewis 1995): when a picture represents a glass that is indeterminately pretty full, it will be indeterminate whether the picture represents a glass that is pretty full. But I wouldn't want to suggest that those who give up UNKNOWABILITY will be forced to think of belief in this way.

[20]For an appeal to introspection along the lines, see Schiffer 2000a, p. 336.

[21]Astronomical texts generally describe Arcturus as an "orange-red" star.

[22]Explanations along these lines are also proposed by Tappenden (1993, p. 15) (who attributes the point to Kripke), and by Braun and Sider (MS, p. 16).

[23]Proposals along these lines are mooted by Fine (1975, p. 139–140), who suggests that the strong disjunction of '<u>P</u>' and '<u>Q</u>' might be defined as 'clearly <u>P</u> $\lor$ clearly <u>Q</u>, and by Field (MSb), who suggests that the weak disjunction of '<u>P</u>' and '<u>Q</u>' might be defined as '$\sim P \rightarrow Q$', where '$\rightarrow$' is a conditional having various desirable properties.

[24]The contextualist also faces a version of this problem. If the logical connectives take on their "strong" senses only in sentences which suggest unanswerable questions, we will never actually be called upon to exercise our understanding of these strong senses in deciding whether to assert a sentence, since we know in advance that sentences which suggest unanswerable questions are always going to turn out to be semantically indeterminate.

[25]Alternatively, the grand convention could be that we are to assert only those sentences that are assigned "true" by <u>at least one</u> of $\mathcal{L}_1, \mathcal{L}_2, \ldots \mathcal{L}_n$. This would mean that sentences which are assigned "true" by only some of $\mathcal{L}_1, \mathcal{L}_2, \ldots \mathcal{L}_n$ are semantically indeterminate not because both they and their negations are forbidden, but because both they and their negations are permitted. If we think that English contains <u>both</u> sorts of semantically indeterminate sentences, the grand convention will need to take a more complicated form.

[26]Lewis 1975 adds conventions of "trust" to these conventions of "truthfulness": if his reasons for doing so are sound, we might want to give find a way to give an equal role to both sorts of conventions in the analyses of determinate truth and semantic indeterminacy.

[27]Cf. McGee and McLaughlin 1995.

[28]It would be somewhat less implausible to hold that the inference from (25) to the implausible conclusion about ignorance would be blocked by the claim that the <u>proposition</u> that the glass is pretty full is neither true nor false. But the disquotation schemas for the truth and falsity of <u>propositions</u> (assuming that there are such things) are even harder to deny than their counterparts for sentences.

[29]I am glossing over the complications which arise when we try to account for the behaviour of the operators inside the scope of quantifiers and other operators. One advantage of the supervaluationistic semantics sketched at the end of the last section is that it lets us give a straightforward theory that covers all sorts of occurrences of 'determinately' and 'it is indeterminate whether': we can say that on any admissible precisification, 'determinately ϕ' expresses the conjunction of all the properties expressed by 'ϕ' on precisifications that are in the extension of 'admissible precisification' in the given precisification.

References

Barnett, David (2000). 'Vagueness-Related Attitudes.' In Sosa (2000), 302–320.

Braun, David and Theodore Sider (MS). 'Vague, So Untrue.' Available online at `http://fas-philosophy.rutgers.edu/~sider/papers/vagueness.pdf`.

Field, Hartry (2000). 'Indeterminacy, Degree of Belief, and Excluded Middle.' <u>Noûs</u> 34: 1–30.

— (MSa). 'No Fact Of The Matter.' Forthcoming in <u>Australasian Journal of Philosophy</u>. Available online at `http://www.nyu.edu/gsas/dept/philo/courses/factual/papers/Field.pdf`.

— (MSb). 'The Semantic Paradoxes and the Paradoxes of Vagueness.' Forthcoming in J.C. Beall and Michael Glanzberg, eds., <u>Liars and Heaps</u>. Available online at `http://www.nyu.edu/gsas/dept/philo/faculty/field/papers/Semantic.pdf`.

Fine, Kit (1975). 'Vagueness, Truth and Logic.' <u>Synthese</u> 30: 265–300. Reprinted in Rosanna Keefe and Peter Smith, eds., <u>Vagueness: A Reader</u>. Cambridge, Mass.: MIT Press, 1996, pp. 119–150.

Grice, H.P. (1975). 'Logic and Conversation.' In Peter Cole and Jerry Morgan (eds.), <u>Syntax and Semantics, Volume 3: Speech Acts</u>, 41–58. New York: Academic Press. Reprinted in Grice 1989: 22–40.

— (1978). 'Further Notes on Logic and Conversation.' In Peter Cole and Jerry Morgan (eds.), <u>Syntax and Semantics, Volume 9: Pragmatics</u>, 113–127. New York: Academic Press. Reprinted in Grice 1989: 41–57.

— (1989). <u>Studies in the Way of Words</u>. Cambridge, Mass.: Harvard University Press.

Hawthorne, John (forthcoming). 'Vagueness and the Mind of God.' Forthcoming in <u>Philosophical Studies</u>. Available online at `http://www.sims.berkeley.edu/~kahern/hawthorne/vaguenessAndMindOfGod.htm`.

King, Jeffrey C. and Jason Stanley (forthcoming). 'Semantics, Pragmatics, and The Role of Semantic Content.' Forthcoming in <u>Semantics vs. Pragmatics</u>, ed. Zoltan Szabo. Available online at `http://www-personal.umich.edu/~jasoncs/semantics-pragmatics.pdf`.

Lewis, David (1969). <u>Convention: A Philosophical Study</u>. Oxford: Blackwell.

— (1975). 'Languages and Language.' In Keith Gunderson (ed.), <u>Language, Mind and Knowledge</u>, volume 7 of <u>Minnesota Studies in the Philosophy of Science</u>, 3–35. Minneapolis: University of Minnesota Press.

— (1993). 'Many, But Almost One.' In Keith Campbell, John Bacon and Lloyd Reinhardt (eds.), <u>Ontology, Causality, and Mind: Essays on the Philosophy of D. M. Armstrong</u>, 23–42. Cambridge: Cambridge University Press.

— (1995). 'Reduction of Mind.' In S. Guttenplan (ed.), The Blackwell Companion to the Philosophy of Mind, 412–431. Oxford: Oxford University Press.

McGee, Vann and Brian McLaughlin (1995). 'Distinctions Without A Difference.' Southern Journal of Philosophy 33 (Supp.): 203–251.

Schiffer, Stephen (1998). 'Two Issues of Vagueness.' The Monist 81: 193–214.

— (1999). 'The Epistemic Theory of Vagueness.' In James Tomberlin (ed.), Philosophical Perspectives 13: Epistemology, 481–503. Oxford: Blackwell.

— (2000a). 'Reply to David Barnett.' In Sosa (2000), 335–343.

— (2000b). 'Vagueness and Partial Belief.' In Sosa (2000), 220–257.

Soames, Scott (1999). Understanding Truth. Oxford: Oxford University Press.

Sorensen, Roy (1988). Blindspots. Oxford: Oxford University Press.

Sosa, Ernest (ed.) (2000). Philosophical Issues 10: Skepticism. Oxford: Blackwell.

Tappenden, Jamie (1993). 'The Liar and Sorites Paradoxes: Towards a Unified Treatment.' Journal of Philosophy 90: 551–577.

Weatherson, Brian (MSa). 'Many Many Problems.' Forthcoming in Philosophical Quarterly. Available online at http://www.brown.edu/Departments/Philosophy/homepages/weatherson/manymany.pdf.

— (MSb). 'Vague Terms with Sharp Boundaries: Reply to Greenough.' Available online at http://www.brown.edu/Departments/Philosophy/homepages/weatherson/cpg.pdf.

Williamson, Timothy (1994). Vagueness. London: Routledge.