

Vagueness Without Ignorance

Cian Dorr


Draft of March 22, 2003.

No comment too large or too small!


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 'Will Bush win the next election?'. 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?' Glasses that are two-thirds full are borderline cases of the predicate 'pretty full'; characteristically, when a is a borderline case of the predicate ' F ', we are motivated to avoid either asserting or denying the sentence ' a is F ' by considerations that have nothing to do with ignorance. In the first three sections of this paper, I will show how this "no-ignorance" view can be developed into an illuminating account of the nature of vagueness and indeterminacy as essentially linguistic phenomena. The remainder of the paper will be spent addressing a powerful objection to the no-ignorance view, due to Timothy Williamson (1994) and other proponents of the epistemic theory of vagueness.

1 A simple vague language

Suppose A and B have instituted a primitive signalling system. They will explore the jungle independently, looking for fruit-bearing trees. When one of them finds such a tree, she will make a noise: either a hoot or a yelp. (A and B 's vocal apparatus doesn't allow them to make any other sounds.) When one of them hears a hoot or a yelp, he will come to believe that the other has found a fruit-bearing tree. Moreover, he will take 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 will favour hypotheses according to which the tree has more fruit; if it was a yelp, he will favour hypotheses according to which the tree has less fruit. Assume for simplicity¹ that B starts off with credence distributed equally over the possible numbers of fruit from 0 to 100, conditional on a fruit-bearing tree being found by A —a state that we represent like this




Then B 's credence distribution after hearing A hoot or yelp will end up looking something like this:



¹... and because of limitations in my ability to draw these diagrams...

This state of affairs could persist even if we assume that A and 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 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 *he* has all these features, and so *ad infinitum*). Of course, if B knew all these facts about A , and also knew exactly what probability A assigned to each hypothesis about how B would update his credences in response to a hoot or a yelp, B would not update his credences in the manner represented above; instead, his credences would end up looking like this:



For the pattern of reactions depicted in the earlier graph to be stable, B must be uncertain exactly how A 's credences are distributed among the different hypotheses about B 's pattern of reactions.²

I boldly assert that the situation I have described is one in which A and B are speaking a very simple two-word language. The hoot and the yelp are clearly not precise words in this language; so if they are meaningful at all, they must be vague.

²And therefore B must also be uncertain exactly how A 's credences are distributed among the different hypotheses about the manner in which B 's credences are distributed among the different hypotheses about the way A 's credences are distributed among the different hypotheses about B 's pattern of reactions; and so *ad infinitum*.


It is characteristic of this language that it gives rise to situations in which the decision whether to hoot or to yelp is a *hard* decision. To make this decision in an ideally rational way, *A* will have to carefully consider how her evidence bears on many different hypotheses about *B*'s credences. It's not quite right to say that the difficulty of this decision has nothing to do with ignorance: the decision would certainly be much easier if *A* knew exactly how *B*'s credences would be affected by each option. But it is hard to see how *A*'s decision would be made any easier if we gave her even more knowledge about the *tree* than she already has.³

2 Why don't we answer 'Yes' or 'No' to borderline questions?

We can bring our imaginary simple language closer to actual vague language by imagining a small change in the situation. In the situation I described, *A* has a decisive motivation (assuming she desires to convey the maximum of relevant information to *B*) either to hoot or to yelp whenever she finds a fruit tree: if she remains silent, *B* will not even know that he has discovered a fruit tree at all. But now suppose that *B* can see for himself whether *A* has found a fruit tree, though he cannot see how much fruit is on it. In this new situation, the old practice will become unstable. *A* must now consider a third option, that of remaining silent. Even in the original situation, there will be some natural temptation to remain silent when one finds a tree with around 60 fruit, since the decision whether to hoot or yelp in such a case is a difficult one, requiring *A* to face difficult questions about the likelihood of different sorts of responses by *B*. Moreover, even if *A* knew that the option of remaining silent would leave *B* with credences distributed just as they would have been had *B* seen *A* finding a tree but been unable to hear her, *A* might be motivated to remain silent if she

³Another moral I would draw from this case is that vagueness is not, contrary to the view of my esteemed colleagues Schiffer (1998) and Field (2000, MS), 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, *A* and *B* 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.

was strongly motivated not to decrease B 's credence in the proposition stating the actual number of fruit on the tree in question. Since B can anticipate that the option of remaining silent will be tempting in these ways to A in the “difficult cases”, B will be inclined to take A 's remaining silent as some evidence that A is actually faced with a difficult case. Since A can anticipate that B will react that way, A will now have a new motivation to remain silent in such cases, in order to avoid misleading B . So the result of our change to the environment will be to institute a new practice, in which B 's credences in case A hoots, yelps or remain silent will look something like this:




A '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 that a certain glass 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 as follows:



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 that a certain glass 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 *risky* 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 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.

The asking and answering of yes/no questions is a sort of game; there is a sort of ultra-legalistic way of playing the game on which the only admissible moves for someone who has been asked a question to make are saying 'yes', saying 'no' and saying 'I don't know'. (Philosophers are especially apt to play the game in this way.) If Questioner is in a mood to play the game in the the ultra-legalistic way, he can continue the exchange in BORDERLINE like this:

Questioner: Is the glass pretty full?

Respondent: Well, it's around two-thirds full.

Questioner: I didn't ask you whether it was around two-thirds full, I asked you whether it was *pretty* full. Please answer the question! Is it pretty full, or isn't it?

Respondent may be tempted at this point to answer 'I don't know' on the grounds that it's the least bad option that's consistent with the rules of the game, on the ultra-legalistic interpretation Questioner seems to have adopted. (And if Questioner can be relied on to anticipate this reasoning, 'I don't know' may turn out not to be so misleading after all.) Nevertheless, I think there is a much better way for Respondent to continue the exchange:

Respondent: I'm not in the mood for silly games! The game of asking and answering yes/no questions is an excellent institution in its proper place: it serves us well, for example, in the courtroom. But right now, it serves no purpose. I recognise that by refusing to answer 'Yes' or 'No' (or 'I don't know') to your question I count as having lost the game, on the ultra-legalistic interpretation—but frankly, I don't care.⁴

One might object that this explanation of our unwillingness to answer 'Yes' or 'No' to borderline questions only applies to situations where we really are concerned to avoid misleading the questioner. 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 *primarily* 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 *as if* she were really trying to convey information to Questioner, and will assess her knowledge

⁴It might be useful to compare this with what you'd say if you were faced with someone who insisted on playing the game of yes/no questions in a *hyper*-legalistic way on which even 'I don't know' isn't admissible:

Questioner: Is the glass pretty full?

Respondent: I don't know.

Questioner: I didn't ask you whether you *knew* that it was pretty full, I asked you whether it *was* pretty full. Please answer the question!

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 very 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 another’s “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.

3 Analysing semantic indeterminacy

It is not only in response to vague (or otherwise 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 what are called “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 his 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 borderline question.

What is the relevant difference between PRAGMATIC and BORDERLINE? The following schematic answer, due 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.⁶ This is not to say that Respondent has an *additional* motivation not to say ‘No’ in

⁵Why 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 *safely* on the right hand side of the road doesn’t count as a convention. This regularity is a consequence of a convention—that we drive on the right—together with a non-conventional regularity—that we 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 the 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, and that these facts are common knowledge.

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

⁶It’s not entirely clear whether we should say that in this case ‘Yes’ and saying ‘No’ are *both* forbidden by the conventions, or that neither answer is forbidden. I’m tentatively inclined to prefer the latter option, 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 very much more 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 is related to the fact, noted by many authors (Lewis 1979, *Kamp*, *Raffman*, Soames 1999, Graff 2000...), that our use of vague expressions is governed by a *rule of accommodation*: if I say of a glass that we can all see to be about two-thirds full that it is pretty full, the use of ‘pretty full’ will temporarily be modified, so that if I then say of another glass that only I can see, and that is in fact two-thirds full, that it is pretty full, no-one will be misled.

The claim that the sentences ‘the glass is pretty full’ and ‘the glass is not pretty full’ can both permissibly be asserted in the context of BORDERLINE does not entail that the sentence

PRAGMATIC in addition to her desire not to mislead Questioner, namely the desire not to violate the conventions of language. Rather, the fact that the conventions are what they are entails that one can typically avoid misleading one's interlocutors only by conforming to the conventions. (Perhaps this is because the latter fact is partly constitutive of the fact that the conventions are what they are.)

If this is the right way to think about the difference between BORDERLINE and PRAGMATIC, we should have all the materials for a general analysis of the notions of semantic determinacy and indeterminacy. The basic idea is that a sentence S is determinately true iff anyone who knew all the relevant facts would be permitted to assert S , and forbidden to assert the denial of S , by the conventions of language use; a sentence is semantically indeterminate iff neither it nor its negation is determinately true.⁷ To make this into an acceptable analysis, we will need to get rid of the weasel words 'all the relevant facts'. This can be done by quantifying: the relevant facts, whatever they are, are *some* facts knowledge of which would suffice for one to be permitted to assert a determinate truth and forbidden from asserting its denial; and any facts which had this feature would have to include all the relevant ones. We end up with analyses of determinacy and indeterminacy that look like this:

C1 S is determinately true in context C for population P iff there is some true proposition Q about C such that asserting S would be permitted, and asserting the denial of S forbidden, to any speaker who was in C and knew⁸ Q , by the

'the glass is pretty full and the glass is not pretty full' can permissibly be asserted in the context. Nor does it entail that one could permissibly assert 'the glass is pretty full' and then 'the glass is not pretty full'. The act of asserting the former sentence might change the context to one in which the assertion of the latter context would be impermissible.

⁷You 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. But you should feel free to read this analysis of 'determinately true' as a stipulative definition: my real target is the metalinguistic predicate 'semantically indeterminate'.

⁸Do I really need to use the notion of *knowledge* here? Perhaps I could get away with

conventions of language use prevailing in P .

C2 S is semantically indeterminate in context C for population P iff neither S nor the denial of S is determinately true in C for P .⁹

This allows for two different ways for a sentence S to be semantically indeterminate: asserting S and asserting its denial could both be permitted, or they could both be forbidden. This seems like a useful distinction. The first sort of indeterminacy is to be found in the language of A and B , the fruit-counters from section 1, and arguably in actual vague languages (see footnote 6 above). The second sort of indeterminacy might 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).

I haven’t been able to think of any clear counterexamples to these analyses, which isn’t to say that I’m very confident that there aren’t any.¹⁰ But I’m not going to press this inquiry any further, since the details don’t matter to anything I’ll have to say in the rest of the paper. The point of going this far was to provide support for the no-ignorance theory by talking of *belief*, or credence above a certain threshold.

⁹These analyses seem to presuppose that S has a denial. What about a sentence in a simple language which doesn’t have an operation of denial? The only sensible reading of the analyses entails that any such sentence is determinately true if there is some truth knowledge of which would suffice for one to be permitted to assert it, and otherwise semantically indeterminate. This is a bit surprising, but on reflection I don’t think it’s that hard to accept. There are various ways one might extend the simple language so as to include a denial of S : one admissible way to do it—not the most natural way, admittedly—would be to make the denial of S be a sentence that can never permissibly be asserted.

I wish I had an analysis of the notion of denial. But one can’t do everything at once.

¹⁰One concern one might have is that the analyses seem to be incompatible with the phenomenon of conventional implicature (Grice 1989): ‘Bush is a president but he is a politician’ is supposed to be determinately true despite the fact that the conventions of language forbid anyone who knows that there’s nothing surprising in a president’s being a politician from asserting it. It’s controversial whether there is any such thing as conventional implicature. If we did want to make room for it, it seems we’d have to somehow draw a distinction between the “central” conventions of language relevant to semantic indeterminacy and the “subsidiary” ones involved in conventional implicature.

showing how it might be developed into a fully-fledged account of the nature of vagueness and indeterminacy.

4 Knowledge-reports in borderline cases

It is time to face a powerful objection to the no-ignorance theory.

Granted, Respondent's difficulty in deciding whether to say 'yes' or 'no' in BORDERLINE is not due to Respondent's ignorance of any *precise* facts. Nevertheless, Respondent is ignorant of *some* relevant fact: she doesn't know whether the glass is pretty full! And this ignorance surely plays a crucial role in generating her practical dilemma: for if she knew that the glass was pretty full, or that it wasn't, she would have good reason to answer 'Yes' or 'No', accordingly. So the dilemma faced by Respondent in BORDERLINE is not, after all, essentially different in character from the dilemma she faces in PRECISE: it is just that whereas in the latter case the difficulty is due to ignorance of precise matters, in the former it is due to ignorance of vague matters.

How should proponents of the no-ignorance theory respond to this argument? In my view, we should reject (refuse to assert) the premise that Respondent doesn't know whether the glass is pretty full. If so, we shall also have to reject at least one of the following claims:

- (1) Respondent doesn't know *that* the glass is full
- (2) Respondent doesn't know *that* the glass isn't full

For it seems obvious that the notion of *knowledge whether* can be defined in terms of *knowledge that*, as follows:

WHETHER x knows whether P iff either x knows that P or x knows that not- P .

Rejecting (1) and/or (2) commits us to rejecting the following general schema, which many (most?) theorists of vagueness have found obvious:

UNKNOWABILITY Necessarily, if it's indeterminate whether P , then no-one knows that P .

Of course, if it's indeterminate whether P , no-one could *determinately* know that P or that not- P . Since knowledge is factive, this could be the case only if it were determinately the case that P or that not- P . The suggestion is that if it's indeterminate whether P , there can be people—such as Respondent—for whom it is *indeterminate* whether they know that P .

I have already pointed out that it would probably be misleading for Respondent to say 'I don't know' in BORDERLINE: this answer is liable to make Questioner think that Respondent can't see the glass very well. This observation doesn't suffice to establish that the sentences 'I don't know that the glass is pretty full' and 'I don't know that the glass isn't pretty full' are semantically indeterminate in the context: Respondent's reluctance to assert them could have some sort of pragmatic explanation. What we now see is that if we want to maintain that the considerations which count against saying 'Yes' or 'No' in BORDERLINE have nothing to do with ignorance, we had better take the unassertability of 'I don't know' at face value, as indicating that the sentence is semantically indeterminate.

There is, nevertheless, something *prima facie* counterintuitive about rejecting UNKNOWABILITY. In the remainder of the paper, I will attempt to overcome this intuitive resistance. In section 5, I will present an argument against UNKNOWABILITY that does not presuppose the no-ignorance theory. In section 6, I will respond to an important argument for UNKNOWABILITY: if we sometimes know the answers to borderline questions, why don't we say? In section 7 I will consider the question what it takes for someone to be a counterexample to UNKNOWABILITY, and consider to what extent ascriptions of other propositional attitudes like belief and desire can be indeterminate in the same way as knowledge-ascriptions. In section 8, I will argue that instances of the law of the excluded middle are determinately true: this will provide a crucial premise for my second independent argument against UNKNOWABILITY, which I will present in section 9.

But is all this really necessary? Isn't there any way for proponents of the no-ignorance theory of vagueness to respond to the objection *without* having to give up UNKNOWABILITY? One step in the objector's argument that we might dispute is the inference from the premise that

(3) Respondent doesn't know whether the glass is pretty full

to the conclusion that

(4) Respondent is ignorant of some fact.

This inference would, I think, be irresistible if we were allowed to appeal to an instance of the law of the excluded middle,

(5) Either the glass is pretty full, or the glass is not pretty full.

From (3) and (5) it follows that

(6) Either the glass is pretty full and Respondent doesn't know that it's pretty full, or the glass isn't pretty full and Respondent doesn't know that it's pretty full.

If we adopt a "deflationary" conception of facts on which 'It's a fact that P ' (or 'The fact that P exists') is equivalent to ' P ', it follows from (6) that there is a fact which Respondent does not know—in other words, a fact of which she is ignorant. But we must adopt such a conception of facts if we aspire to cast any genuine light on the phenomenon of vagueness. If we help ourselves to a "robust" notion of facts on which the equivalence between ' P ' and 'It's a fact that P ' sometimes fails, we will have in effect helped ourselves to the crucial notion of precision or determinacy which we are trying to explain.¹¹

So, to resist the step from (3) to (4), one must reject excluded middle. I will argue against this move in section 8. But I have two reasons for thinking that even if we do reject excluded middle, we will not yet have a satisfactory response to the objection unless we also reject UNKNOWABILITY.

¹¹See Field MS, pp. 1–3.

First: even if we can avoid having to claim that the difficulty of answering borderline questions is due to *ignorance*, we cannot avoid the claim that the difficulty is due to *lack of knowledge* of some proposition: and while this may not be in conflict with the letter of the no-ignorance theory, it seems to be enough to undermine the explanatory power of the picture of vagueness I presented in sections 1–3. According to that picture, the practical dilemmas faced by the fruit-counters *A* and *B* aren't due to any lack of knowledge or certainty on their part (except perhaps knowledge of one another's credences), and our situation when we are asked borderline questions is essentially just a more complex version of theirs. If we accepted that lack of knowledge plays a crucial role in our unwillingness to give straightforward answers to borderline questions, it seems that we would have to concede that the real essence of vagueness is given not by anything along the lines of the analyses of section 3, but in whatever it is that explains why lack of knowledge needn't always involve ignorance.

Second: even if we should *reject* some instances of excluded middle, surely we shouldn't *deny* any: for any sentence of the form $\sim(P \vee \sim P)$ is equivalent, by one of De Morgan's laws, to the explicit contradiction $\sim P \wedge \sim \sim P$. Thus, if we accept UNKNOWABILITY, we will not be able to *deny* (6) or (4) even if we reject excluded middle. So we still won't be able to *assert* the no-ignorance theory, since we won't be able to assert that someone's failure to answer 'Yes' or 'No' to a borderline question has nothing to do with ignorance.¹²

¹²This point can be made more precise as an objection to C1 and C2, my proposed analyses of determinate truth and semantic indeterminacy. According to those analyses, the claim that a sentence '*S*' is semantically indeterminate entails that there is no proposition *Q* such that (i) asserting '*S*' would be permissible, and denying '*S*' forbidden, to anyone who knew *Q*, and (ii) *Q* is true. So in particular, the proposition *that S* cannot be such a proposition. But if we accept UNKNOWABILITY, we will think that the proposition that *S* does satisfy (i). If someone knew that *S*, it would be determinately the case that *S*, so '*S*' would be determinately true, so it would be permissible to assert *S* and impermissible to deny it. So if we are to maintain the analysis, we must claim that the proposition that *S* fails to satisfy (ii): it is not true. And surely, from the claim that the proposition that *S* is not true, it follows that not-*S*. But this is bad news: it entails that anyone who asserts that '*S*' is semantically indeterminate is thereby committed to asserting 'Not *S*'.

5 Omniscience and indeterminacy

In this section I will present the first of my independent arguments against UNKNOWABILITY. Unlike my second argument, to be presented in section 9, this argument will not involve any appeal to excluded middle or other controversial rules of classical logic. It will turn on considerations about omniscience. ‘Omniscient’ means ‘all-knowing’. All of what? All the *facts*, presumably. Given a deflationary conception of facts, it follows that the instances of the following schema are analytic truths:

OMNISCIENCE If x is omniscient, then if P , x knows that P .

Using OMNISCIENCE, modus ponens, modus tollens, conjunction elimination, and conjunction introduction, we can derive a contradiction from any claim of the form

(7) x is omniscient and x does not know that P and x does not know that not- P .

Given UNKNOWABILITY, this would allow us to derive a contradiction from any claim of the form

(8) x is omniscient and it is indeterminate whether P .

But if indeterminacy exists at all, some indeterminacy is necessary: for example,

(9) Necessarily, it is indeterminate whether it is possible for a glass that is two-thirds full to be pretty full.

Hence, the proponent of UNKNOWABILITY must conclude that it is impossible for there to be any omniscient beings.

The idea that considerations about vagueness could justify us in rejecting the possibility of omniscience strikes me as highly unattractive: it seems to make the phenomenon of vagueness altogether too “metaphysically deep”. Perhaps this can be spelled out as follows: angelic beings who spoke a perfectly precise language should be at least as well placed as we

are to engage in metaphysical inquiry, such as inquiry into the question whether there are any omniscient beings.¹³ But such beings wouldn't even be able to formulate this argument from vagueness to the impossibility of omniscience in their language. Nothing they might discover in their anthropological investigations of us speakers of vague languages would strike them as having any bearing on their theological inquiries.

Williamson writes that

[T]o repudiate the very possibility of omniscient speakers. . . is to endorse a strong form of the view that vagueness is an epistemic phenomenon, for it is to treat ignorance as an essential feature of borderline cases. (1994, p. 201)

This seems a bit too strong: one could hold that vagueness is incompatible with omniscience without agreeing with the epistemicists that vagueness can be *analysed* in terms of ignorance. Nevertheless, I think Williamson is right at least to this extent: the position of someone who takes vagueness to be incompatible with omniscience—even if they reject excluded middle—shares many of the features that make epistemicism so deeply implausible.

The proponent of UNKNOWABILITY might attempt to resist this argument by denying OMNISCIENCE. For example, John Hawthorne (MS) advocates replacing OMNISCIENCE with the weaker claim that an omniscient being knows that P whenever *determinately* P . This strikes me as a cheat, which is just to say that OMNISCIENCE seems to me to be an obvious analytic truth. But there are other reasons for finding this option unattractive. We have to be careful: no-one who denies excluded middle will accept the schema

(10) Necessarily, any omniscient being knows whether P

since, given the equivalence WHETHER, this entails that if there are any omniscient beings, excluded middle holds. But even if we don't want to *assert* all the instances of (10), it seems

¹³This claim could be resisted: one might, for example, claim that speakers of a precise language couldn't even grasp the property expressed by our word 'omniscience': what they mean by 'omniscient' is what we mean by 'omniscient about precise matters'. But it is not at all clear that it is possible to be omniscient about precise matters without being omniscient *tout clear*.

very strange to assert the *negation* of any of them—i.e. to assert something of the form

(11) Possibly, there is an omniscient being who doesn't know whether *P*.

To my ear at least, this just sounds awful. But deniers of OMNISCIENCE who accept UNKNOWABILITY, WHETHER and the possibility of omniscience must assert a great many sentences of this form: for example,

(12) Possibly, there is an omniscient being who doesn't know whether it is possible for a glass that is two-thirds full to be pretty full.

In the face of this argument, the proponent of UNKNOWABILITY may retreat to a weaker claim, along the lines of

LIMITED IGNORANCE Necessarily, if it's indeterminate whether *P*, then no *normal human being* knows that *P*.

But this position, even more than the rejection of the possibility of omniscience, is tainted with the implausibility of epistemicism: it entails that there is an important difference between my epistemic situation vis-à-vis the question whether this glass is pretty full and the epistemic situation of certain possible superhuman beings. Whereas I determinately can't know that the glass is pretty full, certain possible beings—such as omniscient beings—don't determinately fail to know that the glass is pretty full. What sort of difference between me and these possible superhuman beings could explain this difference in our epistemic situations? It seems clear that if my current level of knowledge regarding the precise level of water in the glass and the sociological facts about the use of the expression 'pretty full' don't suffice for it to be the case that I don't determinately fail to know that the glass is pretty full, a miraculous increase in my capacity for knowledge of these matters isn't going to help. And these are the only precise facts that could conceivably be relevant; so it seems I wouldn't be any better off even if I were omniscient about all precise matters. But it is quite mysterious what more we would have to do to confer total omniscience on a being who was already omniscient about precise matters (and who had the relevant vague concepts).

6 Co-operativeness and ignorance

Why do so many people find it so obvious that we don't know the answers to borderline questions? My guess is that at least some of this feeling of obviousness is explained by the appeal of the following argument:

If Respondent knew whether the glass was pretty full, why wouldn't she say which it is? Being a co-operative, competent English speaker, she would say 'Yes' in answer to the question 'Is the glass pretty full?' if she knew that the glass in question was pretty full, and she would 'No' if she knew that it wasn't pretty full. Since she doesn't in fact say 'Yes' or 'No', by modus tollens she mustn't know whether the glass is pretty full.

This objection seems to depend on the following general principle:

CO-OPERATIVENESS If x is a co-operative, competent English speaker, and x is asked 'Is it the case that P ?', x will say 'Yes' if she knows that P , and 'No' if she knows that not- P .

In my view, this principle should be rejected. A co-operative person is one who is guided by the desire not to mislead her interlocutors; I have already explained (in section 2) why this desire would lead Respondent to choose an option other than saying 'Yes' or 'No' in BORDERLINE. I don't see how adding the assumption that Respondent knows that the glass is pretty full, or the assumption that she knows that the glass is not pretty full, would do anything to undermine this explanation. Certainly *one* effect of Respondent's saying 'Yes' or 'No' would be that Questioner would come to believe that the glass is pretty full, or that it isn't pretty full: if Respondent's only desire was to make Questioner believe whichever of these propositions is true, she would have a decisive motivation to say 'Yes' or 'No'.¹⁴ But

¹⁴If it's indeterminate whether Respondent knows that the glass is pretty full and indeterminate whether she knows that it isn't pretty full, it can't *determinately* be the case that her

in fact, this is not Respondent's only motivation. She wants to avoid misleading Questioner in any way, and that's why she doesn't say 'Yes' or 'No'.

In the remainder of this section, I will present two independent arguments against CO-OPERATIVENESS. The first argument involves another thought experiment involving an omniscient being. Suppose we present God—a co-operative, competent, omniscient English-speaker—with a glass that is two-thirds full, and ask 'Is this glass pretty full?' Assume CO-OPERATIVENESS, and suppose God doesn't say 'Yes'. By CO-OPERATIVENESS, God doesn't know that the glass is pretty full; so by OMNISCIENCE, the glass is not pretty full; so by OMNISCIENCE again, God knows that the glass is not pretty full, so by CO-OPERATIVENESS God will say 'No'.

This is a disastrous consequence. One way to bring this out is to imagine (following Williamson 1994, p. ??) that we have a roomful of co-operative, competent, omniscient English-speakers to whom we ask the question simultaneously. Using CO-OPERATIVENESS and OMNISCIENCE, we can easily show that they all give the same answer. (To be precise: if any of them doesn't say 'No', all will say 'Yes'; if any of them doesn't say 'Yes', all will say 'No'.) So if we present them with a succession of glasses each slightly less full than its predecessor, there will be a point at which they all suddenly switch from saying 'Yes' to saying 'No': a hidden boundary of exactly the sort that seems most objectionable to opponents of epistemicism.

Another way to bring it out is to introduce a 'determinately' operator. Suppose both CO-OPERATIVENESS and OMNISCIENCE are determinately true, and suppose that God is determinately an co-operative, competent, omniscient English-speaker. Using some minimal modal logic we can then bring the argument of two paragraphs back within the scope of the only desire is to make Questioner believe that the glass is pretty full if it is, and believe that the glass isn't pretty full if it isn't. For if this determinately was her only desire, it would be indeterminate whether she has a decisive motivation to say 'Yes' and also indeterminate whether she has a decisive motivation to say 'No'. But since she is determinately rational, she determinately does whatever she has a decisive motivation to do, so it would have to be indeterminate whether she says 'Yes' and simulataneously indeterminate whether she says 'No'. I argue below that this is impossible.

‘determinately’ operator, reaching the following conclusion:

- (13) Determinately (if the glass is pretty full, God says ‘Yes’, and if the glass is not pretty full, God says ‘No’)

This entails that it’s not indeterminate whether the glass is pretty full, unless it’s indeterminate whether God says ‘Yes’ and indeterminate whether God says ‘No’. But how could this be? God can make noises like ‘Nyes’ and ‘Yo’, but it’s more natural to count sounds like these as determinately being neither utterances of the word ‘Yes’ nor utterances of the word ‘No’. Moreover, even if we do think there could be a sound which was both indeterminately a ‘Yes’ and indeterminately a ‘No’, it’s hard to see how this could help. There may be some cases where a conditional with an indeterminate antecedent and an indeterminate consequent is determinately true: examples might include sentences like ‘if the glass is pretty full, the glass is pretty full’, or ‘if glass A is pretty full and glass B is just as full as glass A, then glass B is pretty full’. These are cases where there is some “penumbral connection” between the meanings of the vague words that feature in the antecedent and the consequent (see Fine 1975). But there aren’t any penumbral connections between the vague predicate ‘is pretty full’ and the vague predicates ‘is an utterance of the word “Yes” ’ and ‘Is an utterance of the word “No” ’.¹⁵

Of course, one could stipulatively define ‘co-operative’ in such a way as to make CO-OPERATIVENESS true by definition. But then the question whether a given omniscient

¹⁵Similar points are made by Hawthorne (MS) and Barnett *cite*. What would happen if we maintained that the required penumbral connections did exist? We’d end up having to posit, whenever it’s indeterminate whether P , a sound which determinately counts as a ‘Yes’ if P and counts a ‘No’ if not- P . So, for example, ‘Nyes’ might be a sound which determinately counts as a ‘Yes’ if the glass is pretty full and a ‘No’ otherwise, while ‘Yo’ determinately counts as a ‘Yes’ if the glass is not pretty full and a ‘No’ otherwise. Unless it’s determinate that P if and only if Q , the sound that is associated in this way with the claim that P must be distinct from the sound that is associated with the claim that Q . The predicates ‘is a “Yes” ’ and ‘is a “No” ’ would, as it were, have to be vague along as many different dimensions as all other vague words combined. The picture seems deeply implausible, although it might become a little less implausible if we thought that speaker’s intentions play a crucial role in determining what words a speaker counts as having uttered.

speaker is co-operative will become just as hard as the question whether a given glass is pretty full. For any omniscient speaker, it will be at best indeterminate whether that speaker is co-operative.

This argument against CO-OPERATIVENESS strikes *me* as rather compelling: but as we saw in section 4, anyone who accepts that we can't know the answers to borderline questions already has a strong motivation to reject either the possibility of omniscience or the schema OMNISCIENCE. Let me turn, then, to my second argument against CO-OPERATIVENESS, which does not depend on either of these premises. Define a *cautious* person as one who never goes beyond what she knows in answering questions, so that the following schema is true by definition:

CAUTION If x is a cautious, competent English speaker, and x is asked 'Is it the case that P ?', x will say 'Yes' only if she knows that P , and 'No' only if she knows that not- P .

Caution, so defined, seems to be something that in some sense we *expect* one another to exemplify. As Williamson (2000) and others have argued, knowledge is a *norm for assertion*. 'You don't really know that' is a reproach; one cannot felicitously assert a sentence of the form ' P and I don't know that P '. This holds for answers to yes/no questions just as much as for any other assertion. We might be tempted to conclude from this that co-operativeness *entails* caution. But this is too strong: a co-operative speaker could falsely believe that she knew that P , and in consequence answer 'Yes' to the question 'Is it the case that P '. It is enough if one *tries* to be cautious. Nevertheless, it is surely possible for a co-operative speaker to *succeed* in being cautious, if she forms no false beliefs about her knowledge.

But this is enough to spell trouble for CO-OPERATIVENESS. Suppose CO-OPERATIVENESS is true: then a cautious, co-operative speaker will say 'Yes' in answer to the question 'Is it the case that P ?' *if and only if* she knows that P . It follows that there is always a determinate answer to the question whether a given cautious, co-operative speaker knows that P , unless there is no determinate answer to the question whether the speaker has said 'Yes'

in answer to the question. Indeed, we can drop this qualification, since even if the speaker produces some sound for which it is indeterminate whether it counts as an utterance of the word ‘Yes’, it is hard to see how it could be determinate that it counts as a ‘Yes’ iff the speaker knows that P : the predicates ‘counts as a “Yes” ’ and ‘knows that P ’ don’t have the penumbral connections that would make this possible (see above). But this is a disastrous conclusion: attributions of knowledge, including attributions of knowledge to cautious, co-operative English speakers, are certainly vague.

Isn’t there anything true in the vicinity of CO-OPERATIVENESS? Well, nothing I have said counts against the following claim:

WEAKENED CO-OPERATIVENESS If x is a co-operative, competent English speaker, and x is asked ‘Is it the case that P ?’, x will say ‘Yes’ if she **determinately** knows that P , and ‘No’ if she **determinately** knows that not- P .

We might try appealing to the truth of WEAKENED CO-OPERATIVENESS to explain away the intuitive appeal of CO-OPERATIVENESS. The thought is that when we evaluate an indicative conditional, we imagine ourselves in a situation in which we would be willing to *assert* the antecedent, and then see whether in such a situation we would be willing to assert the consequent. But we would typically be willing to assert the antecedent of CO-OPERATIVENESS only if it were determinately true. Hence, if WEAKENED CO-OPERATIVENESS is true, our standard method of evaluating indicative conditionals will lead to our finding CO-OPERATIVENESS assertable.¹⁶

¹⁶This story about indicative conditionals is a double-edged sword: my opponent who rejects OMNISCIENCE in favour of the schema

If x is omniscient, and determinately P , then x knows that P .

could use it to explain why we find OMNISCIENCE intuitively appealing. But being able to explain why we might be prone to find something appealing even if it were false is not the same as having an argument against it!

7 Indeterminacy in knowledge and other propositional attitudes

If it's not determinately the case that Respondent doesn't know whether the glass is pretty full, counterexamples to UNKNOWABILITY—cases where it's indeterminate whether P and also indeterminate whether a certain person knows that P —must be quite common. What does it take to make oneself a counterexample to UNKNOWABILITY? The case of Respondent suggests that for many substitutions for P , it is sufficient if one knows as much about the precise facts upon which the question whether P supervenes as any normal human being could know, has a normal grasp of the meaning of the English sentence ' P ', and meets a certain threshold of rationality and reflectiveness.¹⁷ But one doesn't have to know as much as this about the underlying precise facts to be a counterexample to UNKNOWABILITY. The following conditional looks determinately true: if Respondent, who knows that the glass is between 60% and 70% full, knows that the glass is pretty full, so does a less opinionated counterpart of Respondent who knows only that the glass is between 60% and 90% full. If so, then since it is indeterminate whether Respondent knows that the glass is pretty full, it is also indeterminate whether her less opinionated counterpart knows that the glass is pretty full.

In the light of this, the following general principle seems quite plausible:

EXPORTATION If x knows that it's not determinately the case that not- P , then it's not determinately the case that x doesn't know that P .

I wouldn't want to suggest, however, that the only way to be a counterexample to UNKNOWABILITY is to satisfy the antecedent of this principle. It might be enough, for example, if one didn't determinately fail to satisfy the antecedent. Also, we might be reluctant to

¹⁷The latter two clauses are in there to allow for the intuitions (i) that Respondent would not know that the glass was pretty full if she mistakenly thought that the expression 'pretty full' was used by English speakers in the way the expression 'pretty empty' is actually used, and (ii) that Respondent would not know that the glass was pretty full if she was intractably disposed to use the expression 'pretty full' the way everyone else uses 'pretty empty', in spite of her better judgement. I am not sure whether I share these intuitions.

attribute knowledge that it's not determinately the case that not- P to people who don't understand, or radically misunderstand, the word 'determinately'; but this need not prevent one from being a counterexample to UNKNOWABILITY.

If God is a determinately omniscient being, then determinately, if the glass is pretty full, he knows that it is pretty full, and if it isn't pretty full, he knows that it isn't. Is Respondent similar in this respect to God? Here is an argument that she is not. Suppose the glass Respondent is looking at is in fact 65.1% full.

- (i) Determinately, Respondent doesn't know any proposition that entails that the glass is at least 65% full.
- (ii) Determinately, the glass is at least 65% full.
- (iii) It's not determinately not the case that: necessarily, the glass is pretty full iff it is at least 65% full.
- (iv) So it's not determinately not the case that: the proposition that the glass is pretty full entails and is entailed by the proposition that the glass is at least 65% full.
- (v) So it's not determinately not the case that: the proposition that the glass is pretty full entails and is entailed by the proposition that the glass is at least 65% full AND the glass is at least 65% full AND Respondent doesn't know any proposition that entails that the glass is at least 65% full.
- (vi) So it's not determinately not the case that: the glass is pretty full and Respondent doesn't know that the glass is pretty full.
- (vii) So it is not determinately the case that if the glass is pretty full, Respondent knows that it is pretty full.

Why believe (i)? I can think of two reasons. First: Respondent doesn't *believe* any proposition that entails that the glass is at least 65% full: her beliefs don't even implicitly commit

her to anything much stronger than the claim that the glass is between 60% and 70% full. Second: even if Respondent does believe some such proposition, it is hard to see how this belief could amount to knowledge, since such a belief would be radically *unsafe*: it would be false even if the glass were only a tiny bit less full than it actually is.

This argument depends on the fact that Respondent's perceptual access to the underlying precise facts is imperfect. But not all vague questions are like that: there are some for which we seem to be in a position to know everything there is to know about the relevant precise facts. For example, it's indeterminate whether Hibben Apartments are in Princeton: that's because it's indeterminate whether 'Princeton' stands for Princeton Borough, or for 'Greater Princeton' which includes both Princeton Borough and Princeton Township.¹⁸ If I know that Hibben Apartments are in Greater Princeton but not in Princeton Borough, it's hard to see what more I could need to stand in the same relation as God to the same relation to the question whether they are in Princeton: i.e. to be such that determinately, if they are in Princeton I know that they are, and if they aren't in Princeton I know that they are.

Knowledge entails belief; so lack of belief entails lack of knowledge; so determinate lack of belief entails determinate lack of knowledge. So in *BORDERLINE*, it can't be determinately the case that Respondent doesn't believe that the glass is pretty full, or that she doesn't believe that the glass isn't pretty full. To make sense of this, we must reject the popular view that belief is something like internalised assertion, or a disposition to engage in internalised assertions. For it is obvious enough that Respondent isn't disposed to *internally assert* that the glass is pretty full, or that it isn't pretty full.¹⁹ It would be better, for our purposes, to think of belief as analogous to a *picture* (cf. Lewis 1995). If a picture represents a glass of which it's indeterminate whether it is pretty full, it will be indeterminate whether that picture represents a glass that is pretty full.

¹⁸*Lewis's example—cite*

¹⁹Perhaps there is a sense of 'belief'—'occurrent belief' or 'conscious belief'—which can correctly be thought of as internalised assertion. But this doesn't seem to be the sense of belief on which it's true to say that knowledge entails belief.

The surprising sort of indeterminacy in belief-ascriptions can arise even when it is not accompanied by any corresponding indeterminacy in knowledge-ascriptions. Suppose that it is indeterminate whether glass g_1 is pretty full, and that A is as well-placed to know whether g_1 is pretty full as any human being could be, so that it's indeterminate whether A believes that g_1 is pretty full, and indeterminate whether A believes that g_1 is not pretty full. And suppose that B 's credences about the level of water in glass g_2 are the same as A 's about g_1 , and that B is fully competent in the use of 'pretty full', and sufficiently rational and thoughtful. Then, surely, it's also indeterminate whether B believes that g_2 is pretty full, and indeterminate whether B believes that g_2 is not pretty full, no matter what the actual level of water in g_2 might be. We might want to go further than this, and claim that in this situation there is a penumbral connection between A 's beliefs and B 's: determinately, if A believes that g_1 is pretty full, B believes that g_2 is pretty full, and that if A believes that g_1 is not pretty full, B believes that g_2 is not pretty full.

Does an analogue of EXPORTATION hold for belief? The only potential counterexamples to such a principle that I can think of involve people who are confused or mistaken about the meaning of the word 'determinately'. Suppose, for example, that Michael falsely believes that 'determinately' is synonymous with 'knowably', and as a result falsely believes that the sentence 'it's not determinately the case that Caesar had an even number of hairs when he died' is true in English. Arguably, this may suffice for him to believe that it's not determinately the case that Caesar had an even number of hairs when he died. But it seems clear that Michael determinately *doesn't* believe that Caesar had an even number of hairs when he died. To avoid this problem, we might want to strengthen the antecedent of the analogue of EXPORTATION for belief to rule out cases where the subject's belief is due in this way to a misunderstanding of the concept of indeterminacy. But I won't attempt to make this precise.

Besides knowledge and belief, what other sorts of propositional attitudes can exhibit the surprising kind of indeterminacy? It seems pretty clear that *desire*, for example, goes


the same way as belief. Suppose that Fred strongly desires that a certain glass should be between 60% and 70% full, and that he is competent in the use of the expression ‘pretty full’, rational and reflective. If he knew that all his desires were fulfilled, it would be indeterminate whether he knew that the glass was pretty full. It seems plausible that in that case it is indeterminate whether Fred desires that the glass should be pretty full.

However, not all “propositional attitude verbs” behave in this way. For example, if a certain utterance *means that* a certain glass is between 60% and 70% full, or means that it is indeterminate whether it is pretty full, this utterance determinately does not mean that the glass is pretty full. The proposition expressed by the utterance is determinately distinct from the proposition that the glass is pretty full. *Perhaps* it’s possible for there to be an utterance which doesn’t determinately fail to mean that P , and doesn’t determinately fail to mean that not- P ,²⁰ but this certainly isn’t the typical case.


8 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 over the spectral colours, the likely result of Respondent’s answering ‘No’ would be a credence-distribution looking something like this:

²⁰For example, suppose I introduce the word *smiracle* into the language by stipulating that ‘ x is a smiracle iff y is a smiracle’ shall be true when, and only when, ‘ x is a circle iff y is a circle’ is true. If this is all I say, perhaps it’s indeterminate whether ‘ x is a smiracle’ means that x is a circle, and indeterminate whether it means that x isn’t a circle.



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 *less* 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:



Is this a genuine case of semantic indeterminacy, or is the difficulty of answering this question due to “merely pragmatic” factors? Theorists of vagueness disagree about this question. Some say that the sentence

(14) The chip is either red or orange

is determinately true but unassertable for pragmatic reasons; others say that it is semantically indeterminate. Those who choose the latter option 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. In particular, they will classify as semantically indeterminate certain instances of the law of excluded middle, like

(15) The chip is either red or not red.

Likewise, if we give a pragmatic explanation of the unassertability of (14), we will almost certainly have to explain the unassertability of (15) in the same way: for surely if (14) is determinately true, so is the claim that if the chip is orange, it is not red; and (15) seems to follow from (14) together with this claim. In this section, I will argue for a pragmatic explanation of the unassertability of sentences like (14), and by extension of instances of the law of excluded middle like (15).

Consider the following arguments:

(16) (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

(17) (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.

(18) (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.

²¹Arcturus is often described as an "orange-red" star.

(16c), (17c) and (18c) seem to be unassertable in the same way as (14). On the other hand, there seem to be many contexts where it would be entirely appropriate, and not at all misleading, to assert (16a), (17a) and (18a). And (16b), (17b) and (18b) 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 some contexts in which, for example, (16a) and (16b) but not (16c) are determinately true.
- (ii) The contextualist option. We could claim that sentences like (16a) and (16c) are in some way ambiguous: on one resolution, both are determinately true; on the other, both are indeterminate. The argument from (16a) and (16b) to (16c) is determinately valid provided the ambiguities are resolved uniformly. But typically, an assertion of (16a) would be interpreted as determinately true, while an assertion of (16c) 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 (16c) 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 (16a)–(18a) are semantically indeterminate, but assertable in many contexts for pragmatic reasons.
- (iv) The liberal option. We could claim that sentences like (16c)–(18c) are determinately true, but unassertable in most contexts for pragmatic reasons.

Before I discuss the merits of these options, I want to spend a little time thinking about the question what explains the difference between our attitudes to the conclusions of these

arguments and our attitudes to their first premises. I tentatively propose the following hypothesis: we don't want to assert 'Arcturus is either red or orange' because this assertion naturally suggests the question 'Which is it, red or orange?', a question which we know we will not be able to answer straightforwardly.²² By contrast, (16a)–(18a) 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, (16a) *will* 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 could imagine a community for whom (14) 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 process. The fact that a sentence naturally carries the mind to a question which it would be impossible to answer straightforwardly 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 2. 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 these complicated calculations when we're deciding whether to assert something like (14). The point is to explain the prevalence of the practice of avoiding asserting such sentences by showing how

²²Cf. Braun and Sider MS, p. 16.

rational considerations could foster the evolution of such a practice.

Let me now present three more pieces of confirming evidence for the “unanswerable question” hypothesis. First, consider two lists of sentences involving definite descriptions:

- (19) (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) Exactly one of the numbers between 10^{27} and 10^{28} is the mass of Jupiter in kg.
- (20) (a) The first human being wasn’t born until long after the dinosaurs were extinct.
(b) The first human being was an 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) Exactly 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 the sentences without misleading one’s audience. This is well accounted for by the hypothesis: each change of wording makes the question ‘Which number *is* the mass of Jupiter in kg?’ or ‘Which creature *was* 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 '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 "unanswerable question" hypothesis: in this context, the suggested question 'Which is it, red or orange?' *does* have a perfectly appropriate straightforward answer, namely 'I don't know'.

Third, many philosophers have noted that instances of the law of non-contradiction, like

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

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

(22) The chip isn't both not red and not orange.

And yet, the only logical rules we need to derive the unassertable (14) and (15) from these sentences are De Morgan's Laws and (in the case of (21)) double negation elimination. This difference is neatly explained by the "unanswerable question" hypothesis: the forms of (21), (22) 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 (16)–(18), we are in a position consider the merits of each of the four possible ways of accommodating the facts in our semantic theorising.

(i) I don't have much to say about the logically revisionary option, according to which the puzzling arguments have determinately true premises and indeterminate conclusions, at least in some contexts. This view rides roughshod over our intuitions about validity. And

intuitions about validity are not to be lightly dismissed: it is largely by means of them that we uncover the compositional rules which explain our ability to understand and use indefinitely many sentences.

(ii) The main 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 be suggested that the only sort of context-sensitivity we need is the familiar phenomenon of quantifier domain restriction: the contexts where

(17a) 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, even if we don't mind the commitment to fuzzy objects. For one thing, it's quite implausible that when we utter

(16a) Any star with a surface temperature between 3000 and 5000 Kelvins is red or orange.

we are restricting the domain to stars to which our colour-words determinately apply. If we were, (16a) could be true in its context even if some stars with a surface temperature between 3000 and 5000 Kelvins were borderline blue-green! Furthermore, I can't see any way for quantifier domain restriction to explain the other phenomena I noted while giving evidence for the “unanswerable question” hypothesis.

Where else could the ambiguity or context-sensitivity be coming from? The most obvious answer attributes it to words 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. Similarly, there will be a strong and a weak sense of ‘some’, of definite descriptions, and perhaps

(if De Morgan’s Laws are to be retained within each context or uniform disambiguation) corresponding senses of ‘and’ and ‘every’.

The main thing to be said against this is that ‘or’ and ‘and’ just don’t *feel* 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 *said that* 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 ambiguity or context-sensitivity. Nevertheless, we should not lightly flout Grice’s maxim that senses should not be multiplied without necessity (Grice 1989).

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

$$\begin{array}{l} p_1 \models q \\ p_2 \models q \\ \hline /p_1 \vee_{\text{weak}} p_2 \models q \end{array}$$

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 interested in improving on natural language. But is it plausible that something like this is the true logic of natural language?

A different version of the contextualist option would trace the ambiguity or context-sensitivity to vague words like ‘red’. It is not so hard to see how this would work for sentences like ‘the chip is red or orange’: we could say that ‘red’ and ‘orange’ have penumbral connections in some contexts (or on some disambiguations), 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 unconnected senses. But how would this approach handle ‘There is a number that is Jupiter’s mass in kg’? Is the idea that ‘Jupiter’ has one sense on which universal instantiation fails? Or is it ‘is the mass in kg of’ that is ambiguous or context-sensitive, despite its apparent precision? Neither option seems very appealing.

(iii) The big thing to be said in favour of the conservative option is that it comports best with our intuitive response to the cases. If I assert (16a), and someone points out to me that it entails (16c), I will most likely react by backpedalling: I’ll say something like ‘I suppose all I *really* 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 be 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 S into another sentence S^* , and proposes that a sentence S is assertable whenever it does not too naturally suggest a unanswerable question, and 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 ‘ P or Q ’ with ‘Not determinately not (P or Q)’, or ‘If not- P then Q ’, and similarly for other connectives. But if anything like this is right—if there is a wide variety of contexts in which ‘ P or Q ’ is treated for communicative purposes just as if it were synonymous with ‘If not- P then Q ’—why wouldn’t it be better to say that in the contexts in question, ‘ P or Q ’ really *is* synonymous with ‘If not- P then Q ’?

The challenge 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 these 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 (15) 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 view is the fact that it doesn’t fit very well with our intuitive reactions to the cases. Phenomenologically, sentences like

(14) 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 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 an overview, see King and Stanley MS). 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 true beliefs. In between, there is a continuum of moderate answers. The liberal view does indeed seem to be inconsistent with the extremely optimistic answer. But the required departure from optimism is not so great. It is *much* 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*, whereas our reaction to (14) 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 view.

Once we can see our way past this problem, there is everything else to be said in favour of this view. 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 conservative option or the contextualist option. It's not that we follow the rule in the sense that we first determine that (14) 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 strikingings seem to 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 (14) assertable, but this disposition is *masked* 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 proposed by the contextualist and conservative views 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.

Define an *admissible precisification* as a maximal classically consistent set of sentences that includes all determinate truths. If the rules of classical logic are determinately truth-preserving, it follows that a sentence is determinately true iff it is a member of all the admissible precisifications. Thus, if we accept not just excluded middle but all of classical logic, the way will be opened for an account of determinate truth in the style of supervaluationism. I am happy to think of the view of vagueness I have been defending as a version of supervaluationism. Indeed, my claims about knowledge and other propositional attitudes can be very naturally accommodated within the supervaluationist framework. The one part of standard supervaluationist dogma that I think we should reject is the identification of truth with supertruth (truth on all precisifications) and the commensurate identification of validity with preservation of supertruth. These doctrines cause all sorts of trouble, and contribute nothing to the supervaluationist's explanations: they deserve to be scrapped.

9 From Unknowability to epistemicism

We have concluded that the 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?

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

(24) There is a least number of hairs that a person could have without being bald, but we will never be able to find out which number it is.

(25) Each person has a mass in grams, but that we can never know what it is, no matter how accurately we weigh them.

Thanks to epistemicists like Williamson (1994), we now have a vivid sense of just what it would be like to seriously maintain such claims. I'm not going to give an argument against epistemicism: my argument is addressed to the large majority of philosophers who share my feeling that epistemicism is utterly incredible.²³ What is it about epistemicism that prompts this reaction? It's certainly not the *analysis* of vagueness in epistemic terms: even if epistemicists dropped the claim to be giving an analysis, claiming nothing more than extensional adequacy on behalf of their account, the view wouldn't look any less implausible. But what more is there to epistemicism, really, than the serious and level-headed assertion of sentences like (23)–(25)?

Maybe it will be suggested that the problem with epistemicists isn't their assertion of things like (23)–(25) *per se*, but their failure to assert some *additional* thing which, if asserted, would make these claims seem less amazing. But what could this additional thing be? Williamson suggests that the notion of ignorance is conceptually tied to the notions of truth

²³This isn't question-begging: an informal survey confirms that there really are a lot of philosophers who find UNKNOWABILITY obvious even though they find epistemicism preposterous. The point of the argument of this section—and for that matter, of the argument about omniscience I presented in section 5—is to show that this position is unstable.

and falsehood, so that one could avoid having to conclude from (23) that human beings are *ignorant* about the question whether the glass is pretty full if one went on to assert

(26) 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 governed by the disquotation schema, so that (26) entails the contradictory sentence

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

I find this argument against (26) completely convincing. But I also don’t agree with Williamson that (26) would let us block the inference from (23) to the conclusion that we are ignorant whether the glass is pretty full, or do anything else to mitigate the implausibility of (23). (26) is a claim about an English sentence. The claim that human beings are necessarily ignorant about the question 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 (23) needs to be supplemented by something like the denial of (26) before it entails anything about our ignorance.²⁴

For similar reasons, I don’t see how the objectionable character of (23) would be blocked if we went on to assert that

(28) The sentence ‘the glass is pretty full’ is semantically indeterminate.

This claim about language use entails that if we *did* know whether the glass was pretty full, we would have some difficulty communicating this knowledge to other English speakers. The natural way to do so, namely asserting ‘the glass is pretty full’, is blocked to us, perhaps because it would make our audience believe certain false propositions as well as the true

²⁴It would be somewhat less implausible to hold that the inference from (23) to the implausible conclusion about ignorance would be blocked by the claim that the *proposition* that the glass is pretty full is neither true nor false. But whatever about the disquotation schema for sentences, the disquotation schema for *propositions* seems to me to be well-nigh incontrovertible (assuming that there are such things as propositions, that is).

proposition that the glass is pretty full. But these observations do nothing to render (23) less incredible: ignorance of facts which could never be communicated in language is, if anything, even harder to accept than other sorts of ignorance.

Would (23) be easier to accept if we went on to assert

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

This raises a difficult question as regards the interpretation of the object-language operator ‘it is indeterminate whether’. My own view is that this operator is best understood as a sort of injection into the object language of the metalinguistic predicate ‘is a semantically indeterminate sentence’.²⁵ It’s not quite clear what this means; but I think that it should at least have the consequence that (29) is no more effective than (28) at making (23) palatable. Other views about the meaning of ‘determinately’ 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 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 psychological 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.²⁶ 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 (23)–(25)? 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 (23)–(25) wouldn’t seem

²⁵I am glossing over the complications which will arise when we try to account for the possibility of quantifying in to contexts governed by ‘determinately’.

²⁶I am thinking here especially of Field 2000.

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 (23)–(25) 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 (23)–(25). 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 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 the epistemic view really is terribly implausible.²⁷ So (23)–(25) must be rejected; and since I have already argued for the law of excluded middle, the only way to do this is to reject UNKNOWABILITY.

²⁷I have mercifully spared you my attempt to argue against epistemicism, which contained nothing that had not already been said better by others.

References

- 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.’ *Noûs* 34: 1–30.
- (MS). ‘No Fact Of The Matter.’ Available online at <http://www.nyu.edu/gsas/dept/philosophy/courses/factual/papers/Field.pdf>.
- Fine, Kit (1975). ‘Vagueness, Truth and Logic.’ *Synthese* 30: 265–300.
- Graff, Delia (2000). ‘Shifting Sands: An Interest-Relative Theory of Vagueness.’ *Philosophical Topics* 28: 45–81.
- Grice, H.P. (1989). ‘Logic and Conversation.’ In *Studies in the Way of Words*, ? Cambridge, Mass.: Harvard University Press.
- Hawthorne, John (MS). ‘Vagueness and the Mind of God.’
- King, Jeffrey C. and Jason Stanley (MS). ‘Semantics, Pragmatics, and The Role of Semantic Content.’ Available online at <http://www-personal.umich.edu/~jasoncs/semantics-pragmatics.pdf>.
- Lewis, David (1969). *Convention: A Philosophical Study*. Oxford: Blackwell.
- (1975). ‘Languages and Language.’ In Keith Gunderson (ed.), *Language, Mind and Knowledge*, volume 7 of *Minnesota Studies in the Philosophy of Science*, 3–35. Minneapolis: University of Minnesota Press. Reprinted in Lewis 1983: 163–188.
- (1979). ‘Scorekeeping in a Language Game.’ *Journal of Philosophical Logic* 8. Reprinted in Lewis 1983: 233–249.
- (1983). *Philosophical Papers*, volume 1. Oxford: Oxford University Press.
- (1995). ‘Reduction of Mind.’ In S. Guttenplan (ed.), *The Blackwell Companion to the Philosophy of Mind*, 412–431. Oxford: Oxford University Press.
- Schiffer, Stephen (1998). ‘Two Issues of Vagueness.’ *The Monist* 81: 193–214.
- Soames, Scott (1999). *Understanding Truth*. Oxford: Oxford University Press.
- Williamson, Timothy (1994). *Vagueness*. London: Routledge.
- (2000). *Knowledge and its Limits*. Oxford: Oxford University Press.