

Indeterminacy, Degree of Belief, and Excluded Middle

1. Referential indeterminacy (for instance, indeterminacy as to what a singular term stands for or what a general term has as its extension) is a widespread phenomenon. Ordinary vagueness is a special case of indeterminacy: for instance, it is indeterminate whether the word ‘rich’ has in its extension certain moderately rich people,¹ and it is indeterminate precisely which atoms are parts of the referent of ‘Clinton’s body’ at a certain time. But there are more interesting examples as well.

One kind of example arises in the context of belief in false theories. Consider the use of the term ‘heavier than’ by pre-Newtonians. Did it stand for the relation of *having greater mass than*, or for the relation of *having greater weight than*? In pre-Newtonian physics there was no distinction between the weight of an object and its mass; and since the term ‘heavier than’ was applied almost exclusively in the context of objects at the surface of the earth where there is an almost perfect correlation between mass and weight, there is little in the pre-Newtonian use of the term that could have settled the matter. Some pre-Newtonian utterances may to some extent favor one interpretation over the other; unfortunately, they were probably about equally matched in importance by utterances favoring the other interpretation. I suppose we could say that the term didn’t stand for either the relation *more weighty than* or the relation *more massive than*. But then what did it stand for? There is no third relation that is a better candidate for what their term stood for; and it seems unsatisfactory to say that it didn’t stand for anything at all, since that would seem to imply that they never said anything true when they said that one thing was heavier than another, which is hard to swallow. The best conclusion seems to be that their term “sort of stood for” each of the two relations, but didn’t *determinately* stand for either. This means for instance that if A is an object on the moon that is more massive than but less weighty than an object B on earth, there is no determinate fact of the matter as to whether A was heavier than B, on their use of ‘heavier than’. (So if we want referential indeterminacy in a singular term, ‘the heavier of A and B’ is an example.) On the other hand, if A is both more massive than and more weighty than B, then we should regard it as determinate that A is heavier than B on their use of ‘heavier than’.

An interesting feature of this example is that just as Newton in effect discovered that ‘heavier than’ was indeterminate between being more massive and being more weighty, Einstein in

effect discovered that ‘is more massive than’ is itself indeterminate, between having more rest mass and having more momentum per unit velocity. And for all we know, future physicists may find distinctions that we miss, giving rise to indeterminacy that we can’t yet be aware of.

Or so it would seem. But Steve Leeds (1997, section IV), while granting that it makes sense to ascribe indeterminacy to terms in earlier theories, has denied that it makes sense to ascribe indeterminacy to our own terms. The underlying rationale for this seems to be a disquotational view of reference, on which reference for own singular terms is pretty much defined by the schema

(R) If b exists then ‘ b ’ refers to b and to nothing else

(and analogously for general terms, etc.); talk of reference for other people’s terms makes sense only relative to a correlation of their terms to ours, and in cases like the pre-Newtonian ‘heavier than’ there is no best translation to use.² But whatever the rationale, the view is *prima facie* surprising, among other things because there are circumstances where it seems quite reasonable to suspect indeterminacy in specific terms in our currently best theory. A possible example: it seems to be generally believed that the various tensor fields that Einstein introduced into gravitational theory make no physical sense on a sufficiently small scale and that the quantum gravitation theory of the future will have to replace them; but as far as I know our best theories of gravitation today still employ them, because we don’t yet have a clear enough idea of what the appropriate replacement terminology might be. This is the sort of circumstance where indeterminacy may be suspected: terms like ‘the Ricci tensor’ are unlikely to straightforwardly refer, but are also unlikely to be straightforwardly denotationless. Of course examples like this depend on the fact that we know specific ways in which our currently best theory is problematic; but if you grant that it makes sense to suspect indeterminacy in those cases, I don’t see why you shouldn’t grant that it makes sense to suppose it in other cases where we have no such specific knowledge. For surely all of us would concede the possibility that our current theory is false, even if (contrary to fact) we had no specific knowledge of its being problematic in certain ways; and examples like ‘heaviness’ and ‘mass’ make clear that often when a theory is false, some of its terms are indeterminate.

A different kind of example comes from Brandom 1996. Imagine that a community of

English speakers was separated from the rest of us prior to the development of the theory of complex numbers, and that they independently developed the theory of complex numbers. However, they developed different symbols than ours for the square roots of -1 : instead of calling them 'i' and '-i', they call them '/' and '\'. (Of course they know that / is -\ and hence that \ is -/, but still they use both symbols.) This presents a problem of translation: which of their two symbols should we equate with our symbol 'i' and which with our '-i'? It seems clear that there is no right answer here, for (unlike 1 and -1) the numbers i and $-i$ are structurally identical. That is, whereas 1 differs structurally from -1 , for instance in being its own square, no such difference distinguishes i and $-i$: if you take any mathematical sentence whatever and substitute '-i' for 'i' in all occurrences, the resulting sentence has the same truth value as the original. Because of this complete mathematical symmetry between i and $-i$, it is hard to see how any possible facts about the mathematical behavior of the other community could give reason for preferring one translation over the other. And it seems clearly incorrect to think that this is purely an epistemological limitation. It isn't that there is a subtle fact as to "the correct translation" that we can never know, it is that there is simply no determinate fact of the matter: the whole idea of a unique "correct translation" is misconceived.

I have presented this example as an example of indeterminacy not of reference but of translation, but (assuming platonism about mathematics, as I shall) we could equally present it as a problem about reference for terms in the other language: there seems to be no determinate fact of the matter as to which square root of -1 the term '/' refers to.³ Of course, their typical utterances (for instance, ' $/^2 = -1$ ') will come out true under either assignment, hence determinately true.

This example is of interest in connection with Leeds' position, because it certainly seems at first glance that if the foreign terms '/' and '\' are referentially indeterminate, then so is our own term 'i': wouldn't it be grossly chauvinistic to suppose that we have the ability to determinately single out one of the square roots of -1 , but no one else can have this ability? The most obvious way to evade the charge of chauvinism is to say that it isn't that our term 'i' is better than their term '/', but simply that our word 'refers' has determinate application in the case of our terms but not in the case of theirs. But that way of avoiding the chauvinism charge is unavailable to anyone

who wants to deny indeterminacy in our own language since it assumes an indeterminacy in our word 'refers'. Moreover, we could modify the example to make it even clearer that there is indeterminacy in our own language, by extending our language to include '/' and '\ ' as well as 'i'. We could now express the indeterminacy at the object level: there is no fact of the matter as to whether $\sqrt{-1} = i$. The obvious explanation of there being no fact of the matter is that both '/' (in this extension of our present language) and 'i' are referentially indeterminate. Again, the conclusion that 'i' is indeterminate doesn't conflict with the determinate truth of our mathematical claims, since they come out true relative to either assignment of a root of -1 to 'i';⁴ only "don't cares" like ' $\sqrt{-1}$ ' have indeterminate truth value.

One worry about allowing indeterminacy in our own language is that this may appear to conflict with acceptance of such "disquotation schemas" as

(R) If b exists then 'b' refers to b and to nothing else

and

(T) For any x, 'P' is true of x if and only if P(x);

and these principles seem central to our understanding of the notions of reference and being true of. A possible response would be to somehow restrict the application of the schemas to determinate language, but I think that this is not entirely appealing. Perhaps, then, the idea of indeterminacy in our own language should be abandoned? But how can it be, without falling into chauvinism? This is essentially the puzzle about indeterminacy that is raised in Quine 1969.

(Quine's answer to the puzzle is a bit obscure.)

Fortunately the choice between abandoning the disquotation schemas and accepting the incoherence of indeterminacy in our own language is unnecessary: there is a very natural account that allows both. The account requires a combination of two ideas.⁵ The first idea is that indeterminacy can be holistic. For instance, in the '/' and '\ ' example, the candidates for the reference of the two terms are correlated: since the speakers accept the principle " $\sqrt{-1} \neq \sqrt{-1}$ ", their practices dictate that neither square root of -1 can simultaneously count as / and as \ , even though their practices don't dictate which of the square roots of -1 to assign to which term. The second idea is that semantic terms like 'refers' can themselves be indeterminate.

How do we combine these ideas in such a way as to make the schema (R) come out true? The idea is simply to suppose that our acceptance of the disquotation schema (R) creates a holistic connection between ‘refers’ (as applied to our own language) and each of the singular terms of our language. In particular, there is a holistic connection between the interpretations of ‘i’ and ‘refers’: any acceptable interpretation that assigns a mathematical object x to the term ‘i’ assigns to ‘refers’ a set that includes the pair $\langle ‘i’, x \rangle$ but doesn’t include the pair $\langle ‘i’, y \rangle$ for any y other than x . Such a holistic connection between ‘refers’ and ‘i’ guarantees that “(If i exists then) ‘i’ refers to i and to nothing else” comes out determinately true, and similarly for all other instances of the disquotation schema (R).

This solution does what we want: it gives a natural account of how “‘i’ refers to i ” can be determinately true even though the apparently analogous claim “‘/’ refers to i ” isn’t. The reason for the asymmetry is that in learning to use the term ‘refers’ we learn to accept (R), and this sets up a connection between the word ‘refers’ as applied to our term ‘i’ and our term ‘i’; it doesn’t set up any connection between ‘refers’ as applied to ‘/’ and ‘i’. We get this asymmetry without chauvinism: our term ‘i’ is just as indeterminate as the foreign term ‘/’, it’s just that the indeterminacy is “hidden” in our ordinary semantic claims because there is a **compensating** indeterminacy in our ordinary semantic vocabulary.

I believe that this account removes one main worry about positing indeterminacy in our own language. But one should not overestimate it. Earlier I mentioned that Leeds’ denial of the possibility of indeterminacy in our own language was probably due to the acceptance of a disquotational view of reference, according to which the notion of reference for our own language is defined by the disquotation schemas. One might think that if that is Leeds’ view, then the account just sketched shows that it is confused:

Of course disquotationalism doesn’t rule out indeterminacy in our terms. If ‘i’ is indeterminate, then when we define disquotational reference in terms of it and the other terms of our language, we get an indeterminate notion of disquotational reference. The fact that it is determinate that ‘i’ refers to i doesn’t show that ‘i’ is determinate, it is compatible with ‘i’ and ‘disquotationally refers’ both being indeterminate.

But it is this imaginary rebuttal of Leeds that would be in error. The part after the first sentence is correct, but it merely shows how a *prior* indeterminacy in ‘i’ would give rise to an indeterminacy in ‘disquotationally refers’. But presumably Leeds argument is that there is no way to make sense of the prior indeterminacy in ‘i’, if we recognize no notion of reference beyond the disquotational.

One way around Leeds’ conclusion would be to deny the premise that reference must be defined in terms of the disquotation schema. But I will try to show that there is another way around his argument, which makes sense of the indeterminacy in our own language independently of the theory of reference and therefore independently of the issue of whether it is defined disquotationally.

Putting disquotationism aside, the idea that we should deny the existence of indeterminacy in our own language would in any case appear almost hopeless: for surely ordinary vagueness is a kind of indeterminacy, and surely vagueness is ubiquitous? But vagueness itself can seem problematic: indeed, there is a central problem about vagueness, much discussed in recent years, that puts Leeds’ worries into sharper focus. I’ll call it “Williamson’s puzzle”, since he has been a main proponent of it (Williamson 1994). There are people who believe this puzzle to be so serious as to cast doubt on whether the phenomenon of vagueness, or of indeterminacy more generally, can be genuine. It will be the subject of the rest of this paper, until the final section when I return to Leeds.

2. Williamson’s puzzle is that for any question whatever, there is a simple and straightforward argument for the conclusion that it has a determinate, objective, factual answer. Applied to an ordinary vagueness case, the argument goes as follows:

1. Joe is rich or Joe is not rich.
- 2a. If Joe is rich, then it is a (determinate, objective) fact that Joe is rich.
- 2b. If Joe is not rich, then it is a (determinate, objective) fact that Joe is not rich.
3. So it is a (determinate, objective) fact that Joe is rich or it is a (determinate, objective) fact that Joe is not rich.

This amounts to saying that there is a determinate, objective fact of the matter as to whether Joe is rich. An analogous argument can of course be given against any claim of the form “there is no determinate, objective fact of the matter as to whether p”; so what we have here generalizes into

an argument that there can be no such thing as referential indeterminacy.

Indeed, the form of argument generalizes to cases that aren't obviously cases of referential indeterminacy. Consider the familiar idea that certain evaluative debates, or certain debates about indicative or subjunctive conditionals, are "nonfactual". One is tempted to say for instance that there is no objective fact of the matter as to which kind of ice cream is better, chocolate or coffee, and that there is no objective fact of the matter as to whether Bizet and Verdi would have been French rather than Italian had they been compatriots. If the argument is right, neither these nor any other claims of nonfactuality make any sense. We have here a very powerful form of argument.

There are philosophers who accept this argument across the board; Williamson himself seems to be one, though as far as I know he has never discussed its implications except in the vagueness case. Suppose that we have enough information about Joe's income, his assets, his liabilities, the economy of his society, and so forth, to be confident that no further such information could help us decide whether he is rich. Even so, Williamson and other "epistemic theorists" hold, there is a fact of the matter as to whether he is rich; it's just that we can never know. Put another way: facts about richness, insofar as they outrun facts about assets, liabilities and the like, are epistemically inaccessible to us, but they are facts nonetheless (facts which an omniscient god would presumably know even though we can't). Similarly, I assume, there is an objective though epistemically inaccessible fact as to whether it was mass or weight that pre-Newtonian uses of 'heaviness' stood for.

To my mind, this position is beyond belief. Epistemic theorists sometimes accuse those who say there is no fact of the matter of being verificationists. But this charge is totally off the mark: in fact, what is wrong with *the epistemic position* is that like verificationism it blurs the important distinction between the unverifiable but factual and the nonfactual. That distinction has been crucial to science: for instance, Lorentz and Einstein agreed that questions about the absolute simultaneity of spacelike separated objects are unverifiable, but disagreed as to whether they were factual; and nearly everyone has assumed this difference in their positions to be substantive. (The difference can be scientifically important: for instance, John Bell (1987) tentatively proposed reviving Lorentz's theory some years back, to provide a more satisfactory interpretation of

quantum mechanics.) Similarly, an important debate in the interpretation of quantum theory has been whether particles have determinate position when a momentum measurement is made; here it is agreed on all sides that if they do, their position at that time is unverifiable. Such disputes aside, there are plenty of examples where our theories dictate that certain intuitively factual questions could never be answered: e.g., certain questions about the details of the interior of a specific black hole in an indeterministic universe.⁶ It seems beyond belief that the question of whether Joe is rich, or whether the pre-Newtonians referred to weight rather than mass, is anything like that.

I have occasionally heard proponents of the epistemic (no-indeterminacy) view of vagueness concede that there is a difference between vagueness cases and the scientific examples, but say that this doesn't go against the epistemic view: they say that the difference between the two sorts of examples is that in the scientific cases it is merely *physically* impossible (or impossible *according to our scientific theories*) for us to find the answer, whereas in the vagueness cases it is *conceptually* impossible for us to do so. But I don't think that the epistemic theory can be defended on this basis.

To see this, we need to ask just what is supposed to be conceptually impossible in the vagueness case. Obviously the claim can't be that it is conceptually impossible for us to know whether a certain person is rich: imagine our discovering that the person has billions hidden in Swiss bank accounts. The three most likely alternatives are

- (A) that it is conceptually impossible for us to know whether a person is rich given that the person is a borderline case of being rich;
- (B) that it is conceptually impossible for us to know whether a given person is rich given that the person's financial situation is ...

(where the blanks are of course to be filled in in such a way that we would intuitively regard anyone for whom those details were true as a borderline case of being rich); and

- (C) that it is conceptually impossible for us to know whether a given person is rich given that the person's financial situation is ... and given that we don't have any way of ascertaining richness except via financial situation *and given that we don't have any way of determining which side of the division between the rich and the non-rich contains financial situation*

...

In each of these, the locution “p is conceptually impossible given q” should be interpreted as meaning that the conjunction p&q is conceptually impossible.

(A) has it that what is conceptually impossible is conjunctions like

(#) Joe is a borderline case of being rich and we know whether he is rich.

But the claim that (#) is conceptually impossible is rather uncontroversial, and of no use to the epistemic theorist: its explanation is (i) that ‘borderline case of p’ just means ‘case that isn’t determinately p or determinately not p’; and (ii) that it is a conceptual requirement on knowledge that one can’t know that p unless it is determinate that p. (The more commonly cited conceptual requirement that one can’t know that p unless p is a special case.) This explanation of the conceptual impossibility of (#) is available to the non-epistemic theorist, for it does not require an epistemological explanation of ‘determinately’; it does nothing to support the epistemic theory.

Another way to put the point is to notice that even someone who thought that you could explain the notion of a borderline case in terms of the *physical* impossibility of knowing would recognize the trivial conceptual impossibility of (#): (#) would hold because of the analysis of borderline case in terms of the physical impossibility of knowing and the trivial conceptual impossibility of

(##) It is *physically* impossible for us to know whether Joe is rich and we know whether he is rich.

(Note that you could replace ‘Joe is rich’ by a statement about the interior of a black hole in (##) without losing the conceptual impossibility, so that the alleged distinction between the two kinds of examples would not arise on this interpretation.)

The case of (C) is similar: sure, it is conceptually impossible that we know that Joe is rich *given the absence of any faculties by which we could find out*, but similarly we have no way of knowing the details of the interior of the black hole given the absence of any faculties by which we could find out (for instance, given that we have no direct faculty of perceiving the interior of the black hole by extra-physical means).

If there is any hope for using the distinction between conceptual and physical impossibility to defend the epistemic theory, it must be by taking the relevant conceptual impossibility of

knowledge to be of the sort (B). But the only way for (B) to be conceptually necessary is for it to be conceptually necessary that we *not* have the faculties mentioned in (C). Could that be a conceptual necessity? Perhaps: maybe it is conceptually necessary that if Joe's financial situation is ... then there is no fact (or no determinate fact) as to whether Joe is rich, in which case we could use the conceptual necessity noted under (A) to argue that it is conceptually impossible to know whether Joe is rich (or to have faculties for knowing it). In other words, certain kinds of *non-*epistemic theorists could hold it conceptually impossible that we have such faculties. But how could an *epistemic theorist* hold this? After all, the epistemic theorist takes the question of whether people in Joe's financial situation are rich to be a matter of determinate fact; the claim that we have no means to detect such a fact must then be viewed as a medical limitation on our part, not a conceptual necessity. Of course, this is a hopelessly unintuitive way to view the limitation, but that is just to say that the epistemic view is hopelessly unintuitive.

3. If what I have said is right, then the initial argument 1-3 has a false conclusion, so it must go wrong somewhere. But where? There are two main options. One is to say that Premise 1 is wrong: instances of the law of excluded middle fail when the disjuncts lack determinate truth value. I'll call this *the no excluded middle option*. The other main response is to keep excluded middle even as applied to vague or indeterminate language, indeed keep classical logic generally for such language, but to reject premise (2): 'it is a determinate fact that p' and 'it is a determinate fact that not p' are genuine strengthening of 'p' and 'not p', and in cases of indeterminacy neither strengthening holds even though excluded middle holds. I'll call this *the classical determinately operator option*.

There is also a third option: to keep both excluded middle and Premise (2), but give up the inference from 'p or q', 'if p then r' and 'if q then r' to 'r'. Giving up that inference makes perfectly good sense if we read 'if p then r' in certain nonstandard ways: for instance, as $Dp \supset r$ or $Dp \supset Dr$, where 'D' means 'it is a determinate fact that' (and 'q \supset s' abbreviates 's or not q'). In that case, the failure of the inference is due to the gap between 'p' and 'Dp': because of that gap, 'if p then r' in the stipulated sense doesn't imply 'p \supset r', so we shouldn't expect the inference to hold. But clearly the rejection of the inference based on such a nonstandard reading of 'if...then' isn't

really a third option, it is just the second option (the gap between ‘p’ and ‘Dp’) in disguise. (I’ll mention another nonstandard reading of ‘if...then’ later, which also allows both excluded middle and Premise (2), and which emerges very naturally in the development of the second option.) It is formally possible to reject the inference even given the standard reading of ‘if p then r’ as ‘r or not p’, but (at least in the context in which excluded middle is accepted) I think that this has little appeal. Consequently, in the rest of the paper I will focus entirely on the first two options. Much of the rest of the paper will be devoted to assessing their respective merits.

There seems initially to be a great deal to be said for keeping classical logic, including the law of excluded middle. For one thing, there are a great many different alternatives to classical logic on which excluded middle is renounced, and if we are to renounce the use of excluded middle in the context of indeterminacy then we will have to decide between them. Second, if we are to give up excluded middle for language that *is* vague or indeterminate, then presumably we should give it up as well for language that we think *might be* vague or indeterminate: we should reason in a way that doesn’t prejudge the issue of determinacy. Of course it is reasonable to suppose that if, while reasoning in a broad logic that doesn’t presuppose excluded middle, we conclude that a certain portion of language is determinate, then this reasoning will license the use of excluded middle in that context. But still, the broader logic that doesn’t presuppose excluded middle will need to be taken as basic whenever there is a serious possibility of vagueness or indeterminacy in the language. A third point is an extension of the second: it is arguable that there are no sentences at all for which serious worries about vagueness and indeterminacy can be excluded; if so, giving up classical logic in the case of vagueness and indeterminacy would seem to require taking the view that *no* instance of excluded middle $A \vee \neg A$ is ever strictly valid (except when A is valid or $\neg A$ is valid), but always requires justification that is to be given in a logic that doesn’t use excluded middle anywhere. And somewhat independent of the third point there is a fourth: if we are to take seriously the idea that vagueness or indeterminacy is a quite widespread phenomenon, then we should consider the possibility that the language in which we discuss the semantics of vague and indeterminate language will itself be vague or indeterminate; and then if classical logic can’t be used with vague or indeterminate language, we won’t even be able to use classical logic in meta-theoretic reasoning about the logic of vague or indeterminate language. I don’t say that any of this

is decisive, but it provides some motivation for taking the classical logic option.

But there are two difficulties with it that must be overcome before the classical determinately-operator option can be deemed fully satisfactory. Before discussing them, let us set aside a verbal issue. Proponents of the determinately operator option differ as to their preferred use of ‘it is true that p’: some take it as equivalent to ‘p’ (call this *the weak reading*), others take it as equivalent to ‘it is a determinate fact that p’ (*the strong reading*). Obviously nothing can hang on whether we use the term ‘true’ in its weak or its strong sense. Perhaps the safest policy is to introduce two distinct words, ‘true_w’ and ‘true_s’, for these notions. Of course, if we use the strong notion of truth, we don’t need the determinately operator *in addition*: ‘it is true that p’ is just another way of saying ‘it is a determinate fact that p’, on the strong reading of true. (So questions about the interpretation of ‘determinately’ can equally be regarded as questions about the interpretation of ‘true’, on the strong reading of ‘true’.)

The first difficulty with the classical determinately-operator option is a completely obvious one: the operator ‘it is a determinate fact that’ (‘determinately’, for short) would seem to require some sort of explanation. (Using the term ‘it is true that’ in place of ‘determinately’ obviously would solve nothing: we’d be invoking a sense of ‘true’ in which ‘it is true that p’ isn’t simply equivalent to ‘p’, so an explanation of this seems in order.) Moreover, the explanation of ‘determinately p’ can’t be anything like ‘p, and we might find out that p’, for that would collapse the determinately operator view into the epistemic view that we have rejected. But then, what is the explanation? It doesn’t seem at all easy to provide. (Of course we can partly explain it, by citing certain laws that it must obey: for instance, stipulating that it obeys the laws of the modal system T, or S4, or S5. But obviously this is not nearly enough to settle its meaning uniquely.)

The most popular way to try to explain ‘determinately’ (or ‘true’ in the strong sense) is the supervaluational approach: ‘determinately p’ holds iff p is true in all admissible interpretations of the language. But which are the legitimate interpretations? The simple-minded view is that there is only one, the one in which ‘rich’ stands for rich things, ‘Clinton’s body’ stands for Clinton’s body, and so forth ; if so, ‘determinately p’ becomes equivalent to ‘p’, and there is no indeterminacy. To avoid this, we must apparently say something like this: J is a legitimate

interpretation of L iff either it is the correct interpretation of L or there is no determinate fact of the matter as to whether it is the correct interpretation of L. But if we say this we need an antecedent grasp of the idea of no determinate fact of the matter in explaining legitimate interpretation, so that when we then use the latter to explain the former we are going in a circle. The circularity can be disguised a bit more than I have done, but I don't see how to eliminate it. For instance, we might say that the legitimate interpretations are those which assign precisifications of predicates in the original language. Here, a precisification of a predicate is any set that contains everything that that predicate determinately applies to and contains nothing that it determinately fails to apply to; things of which the predicate is indeterminate (that is, of which it neither determinately applies to nor determinately fails to apply to) will be in some precisifications but not others. I've just explained 'precisification' in terms of 'determinate', and no other explanation is obvious, so the explanation of 'determinate' in terms of 'precisification' (via the intermediate notion of 'legitimate interpretation') is not all that helpful in the end.

That's the first worry about the attempt to posit indeterminacy while keeping classical logic. I think it is a fairly serious one. Indeed, I believe that if one supposes that the only way to resolve the first worry is to provide a reductive explanation of determinateness, one will have to conclude that the first worry is totally irresolvable. But I think that the demand for a reductive explanation is unreasonable: after all, we can't give a reductive account of negation, but that doesn't mean we don't thoroughly understand it. What we ought to want, I think, is an account of the conceptual role of the notion which in some loose sense that I will not try to make precise "fixes its meaning close to uniquely". (I say "close to" because we ought to allow that the notion itself be indeterminate.) But the first problem has certainly not disappeared: it is not at all clear how to give the needed specification of the conceptual role. The modal laws that govern this operator are far from fixing the sense uniquely, and it seems unclear what to add to them. (I should add that what we must explain isn't just the conceptual role of assertions of form 'it is determinate that p', but of other constructions in which 'determinately p' is embedded: most notably, of 'it is not determinate that p'.) Let's leave the matter here for now, and go on to the second worry.

The second worry is that even if we had a clear understanding of just how ‘determinately p’ is supposed to strengthen ‘p’, it is not at all clear that we would be done: one could still raise the question of how the purported fact that it is neither determinately the case that p nor determinately the case that not-p is supposed to show why it’s misguided to even speculate whether or not p. And this problem seems especially acute given that by classical logic (which the view assumes) either p or not-p.

One idea for trying to address both of the worries at once (suggested in Field 1994b) is to simply postulate that part of the conceptual role of ‘determinately’ is that we regard it as misguided to speculate about questions that we take to have no determinate answers. As it stands, this seems feeble. One could try to disguise the feebleness by using the locution ‘true that’ in place of ‘determinately’ (in the strong sense, on which ‘it is true that p’ is not equivalent to ‘p’). In this terminology, the idea would be that though either p or not p, still it is neither true that p nor true that not p, and what we are postulating is that it is misguided to speculate about questions that have no true answers. Although the postulate may sound better when put this way, I think that it can’t be: if the postulate is feeble with one terminology it is feeble with the other. (Presumably the view sounds better when put this way only because the analogous claim about weak truth is so uncontentious.)

Why exactly does it seem feeble to simply build into the conceptual role of ‘determinately’ that it is misguided to speculate about questions that we take to have no determinate answers? The reason, I think, is that neither of the original worries seems fully answered. With regard to the first worry, the problem is that unless more is said about the *way* in which we regard it as misguided to speculate about whether Joe is rich, we don’t capture the sense in which this question is indeterminate. (There is a sense in which we may regard it as misguided to speculate about questions whose answers can clearly never be discovered, but if that were the sense in question we would be back in the epistemic view. In another sense, we may regard it as misguided to speculate about questions that people will laugh at us for asking, but that would be even worse at capturing the sense of indeterminacy we want.) It may well seem that the only hope for explaining the relevant sense of “misguidedness” is by saying “misguided because any answer is going beyond the

determinate facts”. Obviously if that were the best we could do then it would be grossly circular to use the misguidedness in an attempt to clarify ‘determinately’.

The second worry doesn’t seem fully answered either. True, if we could fill out our response to the first worry along the lines suggested we would have built into the idea of indeterminacy that it is misguided to speculate on the answers to questions whose answers are indeterminate. But we would not have addressed the worry that this sits ill with the acceptance of classical logic (in particular excluded middle, especially in conjunction with non-contradiction). Doesn’t accepting that either Jones is bald or Jones is not bald, but not both, somehow pull the carpet out of the view that in the intended sense it is misguided to speculate whether Jones is bald?

Despite these doubts, I think there is some hope for the idea of building into the conceptual role of ‘determinately’ that it is misguided to speculate about the answers to questions one regards as having no determinate answers. I will make an effort to do so in Sections 5 and 6, but first I want to say something about the extent to which our problems would be lessened were we to abandon excluded middle for language that is or might be vague or indeterminate.

4. A full discussion of the “no excluded middle option” would be difficult, because there are many different non-classical logics in which excluded middle is abandoned as a general principle, and different considerations apply to different ones. But let us make one important division: between logics in which it is possible to *reject* certain instances of excluded middle without hopeless inconsistency, and those in which that is not so but nonetheless certain instances are *not assertable*. I’ll call the former “radically nonclassical” and the latter “moderately nonclassical” (though these labels could be misleading in a number of ways).

In one respect the radically non-classical logics are more natural in dealing with vagueness: they give a neater explanation of the idea of indeterminacy (or of a borderline case). To say that Joe is a borderline case of richness is simply to say that it is not the case that either Joe is rich or Joe is not rich. But there is a cost: denying instances of excluded middle in this way requires a fairly radical revision of logic: either one that disallows the inference from not-(p or q) to not-p and/or the corresponding inference to not-q, or one that allows the simultaneous assertion both of r and of not-r. (Reason: from not-(p or not-p), the inferences give us not-p and not-not-p; take r

to be not-p.)

I don't say that such a radical revision of logic is out of the question. The first of the two subpossibilities, in the strong form of denying the inference from not-(p or q) *both* to not-p *and* to not-q, seems initially quite natural in the context of vagueness: 'A or B' can with some plausibility be taken to mean something like 'it is either determinate that A or it is determinate that B', in which case it is clear that the inference from not-(p or q) to not-p (and also to not-q) *should* fail. But if we literally propose this as a reading of 'or', the view is only terminologically different from the classical determinately operator view: it is in effect simply the proposal that instead of taking 'A or B' as meaning $\neg(\neg A \ \& \ \neg B)$, we should take it as meaning $\neg(\neg DA \ \& \ \neg DB)$, where D is the determinately operator. (Note that it will do no good to then kick away the ladder of the determinately operator, and simply use 'or' *as if* it were defined in this way: for we could reintroduce 'D' by defining DA as "A or A", and we could reintroduce classical disjunction as well by defining it in terms of negation and conjunction.) Clearly a genuine alternative to classical logic requires messing with 'not' or 'and' (or both) as well as with 'or'. Once these points are appreciated, the first of the two subpossibilities mentioned in the last paragraph looks far less promising, and the best hope (for the radical nonclassicist) would seem to be with allowing the simultaneous assertion of both p and not-p. As much recent discussion has shown (e.g. Priest 1998), it is possible to develop interesting logics in which asserting both p and not-p (and even, asserting their conjunction) is not "hopelessly inconsistent", that is, where this doesn't imply everything. As far as I am aware however, no such "paraconsistent logic" has found very useful application in connection with vagueness. There is certainly more to be said here, but I will not pursue the matter further in this paper.

Turning to the moderate views, how do they explain the idea of Joe being a borderline case of richness? A moderate view will have it that while we shouldn't accept the disjunction 'Joe is rich or Joe is not rich', we shouldn't deny it either; and it is likely to explain the inappropriateness of our asserting it by saying that it isn't true. But for this to make sense, 'true' must mean something like 'determinately true': if it meant 'true' in the classical sense in which "'p' is true" is equivalent to "p", then we couldn't deny the truth of the disjunction without denying the

disjunction itself. Similarly, the view will likely hold that we should deny the disjunction “‘Joe is rich’ is true or false”, where ‘false’ is taken to mean ‘has a true negation’. Again this shows that for either ‘Joe is rich’ or its negation or both, the attribution of truth to the claim is not being regarded as equivalent to the claim itself; ‘it is true that’ is just a determinately operator under another name. If renaming ‘it is determinate that’ as ‘it is true that’ solved the philosophical problems, we could have done that in the classical case as well.

So if the goal is to avoid a special nonclassical notion of truth that amounts to a notion of determinate truth, each instance of excluded middle must go hand in hand with the claim that that instance is true, and also with the claim that one of its disjuncts is true: either we don’t deny any of the three or we deny them all. But the moderate theorist can’t deny all three: not denying excluded middle was what made him moderate. And it seems wholly unsatisfactory to say that in cases of indeterminacy we don’t deny any of the three: after all, the only obvious way to assert that the example is an example of indeterminacy, without using a primitive notion of determinateness, is to assert that neither it nor its negation is true. (A person’s unwillingness to assert ‘Joe is rich or Joe is not rich’ wouldn’t convey to us that he regards ‘Joe is rich’ as indeterminate: it would convey only that he doesn’t know that ‘Joe is rich’ is determinate.)

The moral is clear: if one gives up excluded middle in merely the moderate way, one has just as much need of a determinately operator as does the (non-epistemicist) advocate of excluded middle.

Should we conclude from this that there is no point to abandoning excluded middle for vagueness without denying instances of excluded middle? That would be too quick. We can safely conclude, I think, that the need to *explain* the determinately operator (or if you prefer, the *true that* operator) is just as great on the moderate no-excluded-middle view as it is on the classical logic view. But perhaps the explanation will be easier? One might think this impossible: moderate non-classical logic is presumably a weakening of classical logic, and if you can’t explain the operator with a stronger logic at your disposal, how could you expect to with only a weaker logic? But that argument overlooks the fact that the determinately operator one needs in the nonclassical case is slightly different than the one that is needed classically: in particular, in the nonclassical case

the operator will obey the law $D(A \vee B)$ iff $DA \vee DB$. If the logic admits the deMorgan laws and double negation elimination and the conditional can be defined away, then this law, together with the corresponding law for conjunction (which holds for the classical operator as well), allows us to put every sentence into an equivalent normal form with the following feature: ‘ \supset ’ doesn’t occur at all, and no occurrence of either ‘ \neg ’ or ‘ D ’ has an occurrence of either ‘ $\&$ ’ or ‘ \vee ’ in its scope. (In other words, if a *generalized atomic sentence* is one built from atomic sentences using only D and \neg , then a *normalized sentence* is one built by conjunction and disjunction from generalized atomic sentences.) Indeed, we can normalize a bit further, by disallowing consecutive occurrences of ‘ \neg ’; and if the analog of the S4 law is assumed, as I will, then we can also disallow consecutive occurrences of ‘ D ’. Given this, we can fully explain assertions and denials of determinateness if we can explain assertions of the form DA , $\neg DA$, $D\neg A$, $\neg D\neg A$; $D\neg DA$, $\neg D\neg DA$, $D\neg D\neg A$, $\neg D\neg D\neg A$, and so forth, where A is atomic. So we have reduced the problem of explaining the notion of determinateness quite considerably.

But we certainly haven’t eliminated it: even if we put the “higher order” sentences aside, there is still the problem of explaining what it means to assert that $\neg D(\text{Joe is rich})$ and $\neg D\neg(\text{Joe is rich})$. (Also, what it means to *conjecture* that the corresponding things *might* be the case for Sam.) And here I think that the moderate⁷ no-excluded-middle theorist is no better off than the classical theorist.

Admittedly, the classical theorist had another problem: the apparent oddity of saying that it is misguided to speculate whether Joe is rich (when we believe Joe to be a borderline case) while at the same time saying that of course by classical logic either he is rich or he isn’t. The no-excluded-middle theorist obviously avoids *this* problem, for he doesn’t hold that either Joe is rich or isn’t. I take this to be a serious motivation for giving up excluded middle in the context of vagueness and indeterminacy.

Still, the problem of explaining assertions and conjectures of determinateness and of indeterminateness remains; and as stressed earlier, the nonclassical logic option has its own costs. So let us now go back to the classical determinately-operator option, and see if we can overcome the problems we have found with it.

5. Let us return to the idea briefly mentioned earlier: that it is part of the conceptual role of ‘determinately’ that it is misguided to speculate about the answers to questions that one regards as having no determinate answers. I argued before that saying just this, without saying anything else, was feeble. What I want to do now is to add to it in a way that will make it less feeble. The approach that I will adopt to doing this is influenced by Schiffer 1998, but my proposal will be very different from his in details.

My approach (like Schiffer’s) will be to look at a model of an idealized epistemic agent, and to explain how such an agent might have different attitudes toward sentences or propositions that he regards as potentially indeterminate than toward sentences or propositions that he regards as determinate but difficult or impossible to answer. (I speak of ‘potentially indeterminate’ rather than ‘indeterminate’ because a sentence like ‘Sam is rich’ needs to be handled carefully even when the agent regards it as merely possible that Sam is a borderline case.) To simplify things, I’ll suppose that the agent’s language L is quantifier-free: it is built up from atomic sentences by the truth functional operators plus the operator D .

What are our idealized epistemic agents to be like? One standard idealized model of an epistemic agent is that provided by the crude Bayesian picture, according to which an idealized agent has point-valued degrees of belief in every sentence of his language and they satisfy the laws of probability.⁸ This makes sense (as a crude idealization) when the agent does not recognize any potential for vagueness or indeterminacy in his own sentences, but I don’t think it obvious that the recognition of the possibility of vagueness or indeterminacy should leave it unaffected. I propose a generalization of it, which will have the Bayesian theory as a special case for sentences that the agent treats as not even potentially indeterminate: more exactly, the Bayesian theory will hold in the sublanguage generated from sentences A such that the agent fully believes “Either it is determinate that A or it is determinate that $\neg A$ ”.⁹

It is both heuristically useful and mathematically simpler to start with a probability function P and construct a function Q from it that need not obey the laws of probability; I will take the resulting Q to be appropriate as a degree of belief function. (I will leave open whether there may be other appropriate degree of belief functions not constructible in this way.) P should be thought

of as simply a fictitious auxiliary used for obtaining Q. (The degrees of belief that Q assigns will be very different from the “vagueness related degrees of partial belief” that Schiffer has proposed.)

It should go almost without saying that in talking about a probability function P in the present context, we should assume that the probability function is constrained in the obvious way by the logic of the operator D, which (for reasons that will emerge) I will take to be S4, or more cautiously, a normal modal logic M that is at least as strong as S4. So more explicitly, I assume

(I) If B is provable in M then $P(B) = 1$.

Also, it is natural to assume

(II) If $P(A)=1$ then $P(DA)=1$.

It follows from the first assumption that if A_1, \dots, A_n entail B in the strong sense that $A_1 \& \dots \& A_n \supset B$ is valid in M, then $P(A_1 \& \dots \& A_n) \leq P(B)$. It follows from the two together that if A_1, \dots, A_n entail B on the weaker notion of entailment in which A entails DA, and if $P(A_1 \& \dots \& A_n) = 1$, then $P(B)=1$.

What features do we want Q to have? First, we will want (I) and (II) to hold for Q as well as P. Second, we will want it to be the case that whenever P takes A not to be potentially indeterminate [that is, when $P(DA) = P(A)$ and $P(D\neg A)=P(\neg A)$], then $Q(A)=P(A)$ and $Q(\neg A)=P(\neg A)$; and conversely, that when P takes A to be potentially indeterminate then either P(A) departs from Q(A), or P($\neg A$) departs from Q($\neg A$), or both. But there is a completely obvious way to achieve these goals: simply suppose that for any sentence B of the language, Q(B) is P(DB).

What are the consequences of this? Let A be ‘Sam is rich’ (or the proposition that it expresses in a given context—see previous note), and suppose that the agent’s probability function P assigns degrees q_+ , q_- , and q_0 respectively to ‘DA’, ‘D $\neg A$ ’, and ‘ $\neg DA \& \neg D\neg A$ ’. Since these three sentences are exclusive and exhaustive, the three numbers must add to 1. But then since $Q(A)=P(DA)$ and $Q(\neg A)=P(D\neg A)$, we will get that $Q(A) + Q(\neg A) = 1 - q_0$. If the agent is fairly confident that Sam is a borderline case, q_0 will be close to 1, and we will have a strong deviation from classical probability. If the agent doesn’t know Sam very well, or knows him to be secretive about his finances, q_0 will be moderate and the deviation from classical probability will be less

extreme. And if for some reason the agent is completely confident that Sam is definitely rich or definitely not rich, even if he doesn't know which, then there will be no deviation from classical probability at all.

But though this view licenses departures from classical probability, it does not license departures from classical logic. For since $A \vee \neg A$ is a classical logical truth, $D(A \vee \neg A)$ is provable in our modal logic, so our probability function will give it value 1. So $Q(A \vee \neg A)$ will also be 1, whatever the three q-values are; that is, $A \vee \neg A$ gets degree of belief 1 even when the degrees of belief in the disjuncts sums to less than 1. The classical law for the probability of disjunctions with mutually exclusive disjuncts does not hold for degrees of belief in the context where vagueness is allowed.

The disjunction of the mutually exclusive claims 'DA', 'D¬A', and '¬DA & ¬D¬A' is also a tautology, so the Q-value of this disjunction is 1; but again the disjuncts needn't sum to 1 (unless the modal logic is strengthened to S5). The S4 law does require that the first two disjuncts get Q-values q_+ and q_- respectively, but the Q-value of '¬DA & ¬D¬A' can be less than q_0 . For in S4, the possibilities in which $\neg DA \ \& \ \neg D\neg A$ hold subdivide as follows:

- (B+) $\neg DA \ \& \ D\neg D\neg A \ \& \ \neg D\neg DA$ ["Borderline positive"]
- (B-) $\neg D\neg A \ \& \ D\neg DA \ \& \ \neg D\neg D\neg A$ ["Borderline negative"]
- (D0) $D\neg DA \ \& \ D\neg D\neg A$ ["Definitely indeterminate"]
- (HI) $\neg D\neg DA \ \& \ \neg D\neg D\neg A$ ["Hopelessly indeterminate"]

(In case (HI) Sam is neither definitely rich, definitely not rich, nor definitely borderline; this seems to me a good description of many situations.) If r_{B+} , r_{B-} , r_{D0} , and r_{HI} are the probabilities that P assigns to these, then they must sum to q_0 . Moreover, each of the four statements is consistent in S4, so each of the r-values can be nonzero. But $Q(\neg DA \ \& \ \neg D\neg A)$ is $P(D(\neg DA \ \& \ \neg D\neg A))$, which by S4 must be $P(D\neg DA \ \& \ D\neg D\neg A)$, i.e. r_{D0} . So the Q-values of our three disjuncts add to $1 - r_{B+} - r_{B-} - r_{HI}$.

Similarly, $Q(\neg DA)$ is $P(D\neg DA)$, which inspection of the above reveals to be $1 - q_+ - r_{B+} - r_{HI}$; so we have

$$Q(\neg DA) \leq 1 - Q(A),$$

with the inequality strict unless both r_{B+} and r_{HI} are 0. So though $Q(DA \vee \neg DA)$ is of course 1,

$Q(DA) + Q(\neg DA)$ is merely $1 - r_{B+} - r_{HI}$. Analogously, $Q(D\neg A) + Q(\neg D\neg A)$ is only $1 - r_{B-} - r_{HI}$.

What this suggests, of course, is that for an agent to *treat A as potentially indeterminate* for him to have degrees of belief in it and its negation that add to less than 1. We can also call this der indeterminacy, and then generalize: an agent eats A as potentially (k+1)st order indeterminate iff he treats both DA and D¬A as potentially kth order indeterminate. It is obvious that potential (k+1)st order indeterminacy requires potential kth order indeterminacy, and easy to see that if there is potential indeterminacy in a sentence B with a chain of k occurrence of ‘D’, each in the scope of the previous member of the chain, then there must be at least kth order indeterminacy in one of the atomic sentences in B. Note carefully that with *merely* first order indeterminacy in A, we have that $Q(\neg DA)$ is just $1 - Q(A)$: the first order indeterminacy comes out in these being less than $Q(\neg A)$. But with second order indeterminacy, we typically have that $Q(\neg DA) < 1 - Q(A)$ and $Q(\neg D\neg A) < 1 - Q(\neg A)$, though we may on occasion have just one of these.

It is trivial to verify that if A_1, \dots, A_n entail B in S4 (in the strong sense mentioned above, corresponding to a derivation where the rule A/DA is applied only to the axioms of T), then $Q(A_1 \& \dots \& A_n) \leq Q(B)$. And the conclusion of the Williamson argument, $DA \vee D\neg A$, has Q-value $1 - q_0$. Consequently, the conjunction of the premises $A \supset DA$ and $\neg A \supset D\neg A$ can have a Q-value no higher than $1 - q_0$; indeed, it is easy to see that the value is exactly $1 - q_0$. (The values of the two conjuncts isn’t determined by the q’s and r’s: we can only say that $Q(A \supset DA)$ is in the interval $[1 - q_0, 1 - q_0 + r_{B+} + r_{HI}]$, and $Q(\neg A \supset D\neg A)$ is in the interval $[1 - q_0, 1 - q_0 + r_{B-} + r_{HI}]$.)

It is instructive to define conditional Q-values in analogy with conditional probability: $Q(B|A) = Q(A\&B)/Q(A)$, provided that $Q(A) \neq 0$; undefined or 1 if $Q(A)=0$. It follows that $Q(B|A)$ is just $P(DB|DA)$. Suppose that we add to the language a “generalized Adams conditional” \rightarrow , allowed only as a main connective, governed by the principle that the degree of belief in $A \rightarrow B$ should always be the conditional degree of belief in B given A (at least when that is defined).¹⁰ Taking conditional and unconditional degree of belief to be represented by conditional and unconditional Q, we get that the degree of belief in $A \rightarrow B$ should be $Q(B|A)$. As is well-known, the ordinary Adams conditional, where degrees of belief are assumed to obey the laws of probability, does much better than the truth functional conditional at capturing ‘if ... then’ in

English, at least if we restrict ourselves to sentences in which ‘if ...then’ appears only as the main connective: consider ‘If I try out for the Yankees I will make the team’.¹¹ So it might seem reasonable to conjecture that the generalized Adams conditional is appropriate for vague sentences. If so, then the degree of belief in $A \rightarrow DA$ should be $Q(DA|A)$ (at least when $Q(A) \neq 0$), which is 1. (Note the contrast between Q and probability functions: for them, the conditional probability is never higher than the probability of the material conditional.) This is noteworthy in connection with the Williamson argument (1)-(3): all the premises of the argument get Q -value 1 if ‘if...then’ is read in accordance with the generalized Adams conditional.¹² Of course, the classical rule of inference used in the argument is not valid on this reading of ‘if...then’.

I might remark that the generalized Adams conditional is also illuminating in the context of the Sorites argument: if we understand the Sorites premises as generalized Adams conditionals they are all highly believable, even for the clearest of borderline cases.

I introduced Q by starting from a classical probability function P , but have suggested that P not be taken seriously: except where it coincides with Q , it plays no role in describing the idealized agent. (Obviously the same Q can be generated from many different P 's, unless $Q(A) + Q(\neg A)$ is 1 for all A .) There are really two distinct parts to this idea: first, that Q is a perfectly legitimate belief function, and second, that P is not a legitimate belief function when it generates a Q distinct from it. The second component of this would be hard to motivate if the process that led from P to Q could be iterated to yield a Q^* distinct from Q , but that is not the case: if we define $Q^*(A)$ as $Q(DA)$, Q^* simply is Q . That is my principal motivation for using S4.

The use of S4 might seem objectionable, on either of two grounds. First, it might be thought odd to have the degree of belief in DA always equal to the degree of belief in A , given that DA is stronger than A . It seems to me, though, that the feeling of oddity goes away once one realizes that the degree of *disbelief* in DA can be higher than the degree of disbelief in A (and typically will be when A is judged indeterminate): that is, $Q(\neg DA)$ can be higher than $Q(\neg A)$, and this is symptomatic of the added strength of DA over A . With classical probability the degree of disbelief in A is determined by the degree of belief in A ; when, as with Q -type degrees of belief, this is not so, we must remember to consider degrees of disbelief as well as degrees of belief when comparing an agent's attitudes toward two propositions. (Without modifying the basic theory we

could modify the notion of degree of belief so that it does determine degree of disbelief: say by using the phrase ‘the degree of belief of A’ to stand not simply for $Q(A)$ but for the pair $\langle Q(A), Q(\neg A) \rangle$, or the interval $[Q(A), 1-Q(\neg A)]$, or some such thing. Then the degrees of belief in DA and A will typically differ.)

Second, the S4 law might be thought inappropriate for higher order vagueness. But that is certainly not obvious: indeed, we have seen above that significant higher order vagueness can be recognized in Q as well as in P, as long as S5 is not assumed. Moreover, as Williamson 1994 notes (p.160), the invocation of a definitely operator that does not obey S4 invites its replacement by an operator D^* that functions as an infinite conjunction of D, DD, DDD and so on; and it is bound to obey S4. (As Williamson points out in a footnote, it won’t obey S5 unless D obeyed the Brouwersche law. And the Brouwersche law seems to me implausible for vagueness and indeterminacy: even without the S4 law it rules Case (HI) to be impossible, and with S4 it rules Cases (B+) and (B-) impossible as well.)

I should add that even acceptance of the S5 law wouldn’t rule out “higher order vagueness” in all possible senses of that term. For instance, stipulating that the logic of ‘determinately’ is S5 would seem to be quite compatible with taking ‘determinately’ to itself be vague, and simply requiring that the admissible sharpenings of the notion of ‘determinately’ be required to obey S5. This thought is expressed in the language of admissible sharpenings, which I have held to be not ultimately illuminating (though not illegitimate either); I will say a bit later how this might be captured in the language of idealized agents. For now let me just say that still another component of the intuitive idea of “higher order vagueness” is that ‘appropriate degree of belief’ is vague. In the current framework, that is connected with the idea that there is no one right way to assign Q-values to sentences. This sort of higher order indeterminacy needn’t be accommodated explicitly in the model of idealized agents on offer, because each legitimate choice of Q-values will correspond to one way of making ‘degree of belief’ precise; if the model is appropriate for each such choice, it is appropriate overall. We could consistently suppose that the S5 law should hold for these legitimate choices of Q-values without denying that there is indeed a multiplicity. But let’s avoid further discussion of these issues by sticking to S4.

If, as I have suggested, the classical probability function P should be viewed as having no

“psychological reality”, but is simply a device for generating the genuine degree of belief function Q , then it would be nice if we could develop the laws governing Q -functions without reference to probability functions. To see how to do this, note that for any sentence B whatever, $Q(DB \vee D\neg B) = Q(B) + Q(\neg B)$; and more generally,

$$(*) \quad Q(DB \vee DC) = Q(B) + Q(C) - Q(B\&C).$$

Proof: $Q(DB \vee DC) = P[D(DB \vee DC)]$. But in S4, $D(DB \vee DC)$ is equivalent to $DB \vee DC$, and by the disjunction law for probability functions, we get $P(DB) + P(DC) - P(DB \& DC)$. But also, $DB \& DC$ is equivalent to $D(B\&C)$, so we get $P(DB) + P(DC) - P[D(B\&C)]$; that is, $Q(B) + Q(C) - Q(B\&C)$.

A consequence is

$$(\#) \quad Q(B\vee C) \geq Q(B) + Q(C) - Q(B\&C);$$

we get an inequality where the classical theory has equality. Actually we can do considerably better than this: the proof of (*) generalizes to a proof of

$$(**) \quad Q(DB_1 \vee DB_n) = \acute{O}_1 Q(B_i) - \acute{O}_{i,j \text{ distinct}} Q(B_i\&B_j) + \acute{O}_{i,j,k \text{ distinct}} Q(B_i\&B_j\&B_k) - \dots \\ \pm Q(B_1\&\dots\&B_n),$$

where the ‘ \pm ’ is + if n is odd and – if n is even.¹³ And this yields

$$(\#\#) \quad Q(B_1 \vee B_n) \geq \acute{O}_1 Q(B_i) - \acute{O}_{i,j \text{ distinct}} Q(B_i\&B_j) + \acute{O}_{i,j,k \text{ distinct}} Q(B_i\&B_j\&B_k) - \dots \\ \pm Q(B_1\&\dots\&B_n),$$

which doesn’t follow from (#) alone. Glenn Shafer (1975) constructed an elegant theory of belief functions on the basis of condition (\#\#) (together with the assignment of 1 to logical truths and the stipulation that the degrees be real numbers in $[0,1]$). Shafer was motivated by considerations not having to do with vagueness, and some of the apparatus he develops don’t have any obvious application to vagueness;¹⁴ still, the above derivation of (\#\#) shows that his basic theory of belief functions can be taken over.

Where are we? We’ve seen that if we start from a classical probability function P on a language L that includes a determinately operator (and where P respects the S4 laws for D), then the function Q_P that we obtain from it must be a Shafer function on L . This means, of course, that its restriction to the D -free part of L must also be a Shafer function. And this suggests that talk of the determinately operator may have been, in a sense, unnecessary: the key to regarding a sentence

as indeterminate is simply adopting degrees of belief (in D-free sentences) that accord with something like the Shafer laws; having attitudes toward sentences involving ‘determinately’ isn’t necessary. Still, a more reflective agent can have attitudes toward sentences that contain ‘determinately’: his degrees of belief in sentences in the language that contains ‘D’ will extend his degree of belief function for the D-free part, and accord with the law that $Q(DA)$ always equals $Q(A)$. And we get the nice result that our two readings of “regards A as determinate” coincide: his degree of belief in $DA \vee D\neg A$ will always equal the sum of his degrees of belief in A and in $\neg A$.¹⁵

I close this section by mentioning three technical questions that I haven’t investigated; their answers could well be obvious. First, can any Shafer function on the D-free part of the language be generalized to a Shafer function on the full language (that meets the constraints on D)? I would hope so, but if not, it implies at most that an agent with an unextendible Q-function on the D-free part of his language needs to revise his degrees of belief a bit if he is to introduce the term ‘definitely’. Even that may be misleading, for I have not ruled out that there be other laws on degree of belief functions beyond the Shafer laws that I have derived. That suggests the second technical question: can any Shafer function on the D-free part of the language be obtained from a probability function on the full language? If not, then there are laws that hold for the special Shafer functions that are so derived that don’t hold for all Shafer functions, and we then need to decide whether to impose all, some, or none of those laws on “degrees of belief”. I would not want to impose any such laws merely on the basis of their being required in order that the degrees of belief “ultimately derive from” a classical probability function, but some such laws might seem independently desirable. A third technical question is, when a Shafer function on the D-free part of L is extendible to a Shafer function on the full L, is this extension unique? I would assume not: I would assume for instance that the Q-values of many sentences of form $DA \supset B$ and $A \supset DB$ when A is D-free won’t be fixed by the Q-values of D-free sentences. This is relevant in connection with the question I raised earlier, of what means would be available for recognizing higher order vagueness if for some reason we were to take D to obey the laws of S5. One way to do this, I would think, would be to suppose that there was a large element of arbitrariness in the extension of the Q-function from the D-free language to the language with D. I will say no more about this, though; higher order vagueness is not my topic, and besides, I see no reason to adopt a

logic stronger than S4.

It's time to review the material of this section. I have tried to give a model of an agent with the feature that despite the agent's adherence to classical logic, he treats sentences that he regards as potentially indeterminate significantly differently from ones that he regards as certainly determinate: the degree of belief in the indeterminacy of a sentence A is measured by the extent to which the probability of it and its negation sum to less than 1. There has been no attempt here to *explain* why the agent adopts this differential attitude between sentences he takes as probably indeterminate and sentences he takes as probably determinate: more on that in the next section. But still, the model seems to be one on which "belief" in sentences that one treats as indeterminate "works differently"—obeys different laws—than belief in sentences that one treats as determinate.

This impression might be thought to be due simply to the crudity of the model for factual belief from which we started. That model, recall, took off from the assumption that when indeterminacy isn't in question, the agent can be taken to assign classical point probabilities to every sentence. But even when indeterminacy is not in question, the assumption of classical point probabilities is unduly restrictive (as Isaac Levi has often emphasized: Levi 1974). A better model would start with a (convex nonempty) set S of probability functions on the language. For each P in S , we could construct a degree of belief function Q_P , in just the way done before; we could then take the set of all Q_P for P in S as what represents the agent's degrees of belief.

At first glance, this complication doesn't alter the important point: sentences viewed as indeterminate will differ in how they function from sentences viewed as determinate. On second glance, this may not be quite so clear. For from the nonempty set S of probability functions used to represent the agent's degrees of belief in determinate sentences, one can generate a numerical degree of belief function G_S : for any sentence A , $G_S(A)$ is taken to be the greatest lower bound of all $P(A)$ for P in S . If we do this, then $G_S(A) + G_S(\neg A)$ will typically be less than 1, and so we will have "degree of belief functions" that look superficially like the Q -functions I have constructed, quite independent of vagueness or indeterminacy. So the worry is that by moving from a single probability function to a set of many probability functions even when indeterminacy is not at issue, we are in effect simply moving to a nonclassical degree of belief function G_S whose nature might mask any further nonclassicality due to indeterminacy.

I can't fully treat the matter here, but there are two reasons to think that this worry is not serious. First, the function G_S constructed from a nonempty convex set S of probability functions will not typically obey the Shafer law (##) or even (#).¹⁶ Second, G_S isn't adequate as a representation of the degrees of belief in determinate propositions anyway, for it gives only partial information about the set S of probability functions from which it was obtained, and the information lost is crucial.¹⁷ These points should do a great deal to allay the worry.

Moreover, even if this worry about the Levi model of idealized agents were to prove correct, I think that a still fuller extension of the model of idealized agents would restore the distinction in function between beliefs that the agent regards as potentially indeterminate and beliefs for which this is not so. Recall the initial intuition: that it is *misguided* to speculate as to whether Joe is rich (where Joe, unlike Sam, is someone I know enough about to be quite sure that he is a borderline case (as I am currently using 'rich')). In the case where there is no indeterminacy, I don't think that precise degrees of belief are misguided, I just don't think that they are mandatory. So a fuller psychological model needs something that evaluates our degrees of belief, and in this way the psychological distinction should be restored.

6. I've gotten into a lot of detail, but now it's time to return to the worries about the determinately operator approach to vagueness and indeterminacy. When we left the matter at the end of Section 3, we had a version of the approach on which it was somehow built into the meaning of 'determinately' that it was misguided to even speculate on questions when the answers were thought indeterminate; or to put it another way, where to regard a question as indeterminate consists in part in our regarding it as misguided to even speculate on the answer to it. The first problem was that the relevant sense of speculation being "misguided", or at least the relevant sense of its being "regarded as misguided", needs specification: without that, the view seems to collapse into something like the epistemic view. Does the kind of idealized model I have discussed help with this problem? I think it does. In terms of the idealized models we don't have any analog of the idea of speculation actually being misguided, but I don't think we need that. What the models do is to build something like the idea of "regarding speculation as misguided" into the core of the account of proper psychological functioning. That is, it is part of the idealized psychology that we simply refuse to adopt classical point valued probabilities in certain cases; this is taken as a natural

state of affairs, not requiring special explanation. In terms of it, we can explain what it is to regard a sentence indeterminate: we regard a sentence indeterminate to the extent that the non-classical probabilities we assign to it and its negation sum to less than 1.

One might object that this explains what it is to *regard* a sentence indeterminate, but that what we should want is an explanation of what it is for a sentence to *be* indeterminate. But let's look at a somewhat similar issue: the notion of objective chance. As is widely recognized (e.g. Lewis 1980, Leeds 1984, Skyrms 1984), it seems impossible to provide a non-circular reductive explanation of the notion of objective chance; all we can do is specify its conceptual role. And the central component of the conceptual role of objective chance is provided by the connection between *opinions about* objective chance and degrees of belief in other matters. Indeed, at least in many important cases, an idealized agent's degrees of belief in the objective chances are completely determined by their degrees of belief in specific outcomes: cf. diFinetti's theorem and its many extensions.¹⁸ I think it is fairly uncontroversial that the meaning of 'chance' is determined by its conceptual connection to degrees of belief in matters other than chance, in roughly this way. Different opinions are possible on what if anything this shows about the metaphysics of chance: one can try to combine it with a "realist" view of chance, or with a view that is in one way or other "anti-realist" (e.g. a view on which attributions of chance are "mere projections" of our degrees of belief about other matters). Some "anti-realists" will hold that talk of degree of *belief* in chances isn't strictly appropriate, since statements of chance are not straightforwardly factual, and that one should speak instead about degree of *acceptance* of chance-sentences; others will take a more relaxed attitude toward the use of 'belief'. I don't propose to enter into these issues, or even to decide whether they are meaningful. But I think that whatever view one takes on those issues as regards objective chance, one can take the same stand as regards determinateness. Here too, we can explain degree of "belief" about determinateness in terms of its connection with degrees of belief in ordinary sentences: at least if we ignore complications arising from the Levi model, we can say that to "believe" that a sentence has a determinate truth value *just is* to have degrees of belief in it and its negation that sum to less than 1. Whether this is adopting an "anti-realist" view of determinateness, and whether it is appropriate to talk of our attitudes toward statements of determinateness as literally "beliefs", need not be settled here.

So I think that the idealized model does a good deal to explain the meaning of ‘determinately’. But what about the second worry: that there seems to be something odd about thinking it misguided to speculate whether or not Joe is rich, given that one accepts that either Joe is rich or Joe is not rich. Here too I think the model helps. First, the model provides a clear picture of epistemic principles that license full belief in excluded middle even while dealing with sentences that are taken to be indeterminate. I think it was largely the absence of such a model that made it seem odd to adhere to excluded middle while recognizing indeterminacy. Second, the feeling that this is odd stemmed in part I think from the assumption that if we assume that Joe is rich or Joe is not rich, it can’t be misguided to assign a classical probability to the claim that Joe is rich. Assigning a classical probability does seem intuitively doubtful in this case; moreover, it might be argued that if we can assign a classical probability to the statement that p then we can speculate as to whether p . (A crude way to argue this would use the supposed connection between classical probabilities and betting behavior, and the fact that it is not misguided to wonder whether one has won one’s bets.) So by rejecting the assumed connection of classical logic with classical probability, the model helps relieve the feeling of oddity in assuming classical logic for indeterminate sentences.

7. One of the features of the account I have given is that it takes attributions of indeterminacy to arise primarily for sentences rather than for terms: terms are to be regarded as indeterminate only because sentences that contain them are to be so regarded. This may seem to raise a problem for one of the examples given in Section 1: I claimed that there was no fact of the matter as to which square root of -1 ‘ i ’ referred to, but we don’t in our language have means of referring to square roots of -1 that aren’t conceptually tied to ‘ i ’, so we don’t have indeterminate sentences, so what can it amount to to say that ‘ i ’ is indeterminate? But this isn’t really a problem: for since we think there are two roots of -1 , we can always introduce new names for them, not linked to ‘ i ’: for instance, ‘ \prime ’ and ‘ \backslash ’. The indeterminacy consists in the fact that if we do that, our degrees of belief in ‘ $\prime=i$ ’ and ‘ $\backslash=i$ ’ should each be 0 , even though the degree of belief in the disjunction should be 1 . (It should be clear that this account of the indeterminacy does not rely at all on the theory of reference: thus it is available to anyone who thinks of reference in one’s own language as defined by the disquotation schema, and thus solves the problem for the disquotationist mentioned near the end of Section 1.)

I now want to apply the ideas of the two previous sections to another example of indeterminacy. As is well known, the size of the set of real numbers is radically unsettled by all the usual axioms of set theory (and by any known way of consistently extending these in a non-ad-hoc fashion): the axioms imply that the set of reals is uncountable, but it is consistent to suppose its size to be \aleph_1 or \aleph_{817} or $\aleph_{\aleph+23}$ or nearly any other uncountable cardinal you like. A possible view is that even though our axioms don't settle the matter (nor do any extension of the axioms that we find at all "evident"), still our set theoretic concepts are perfectly precise, so that there is an objectively correct answer that we will probably never know. But an alternative view, which I find far more plausible, is that our set-theoretic concepts are indeterminate: we can adopt any one of the above claims about the size of the continuum we choose, without danger of error, for our prior set-theoretic concepts aren't determinate enough to rule the answer out.¹⁹ Prior to our choice there is no determinate fact of the matter. In my view, this does not require that we give up classical logic when doing set theory. As I once put it,

...standard mathematical reasoning can go unchanged when indeterminacy in mathematics is recognized: all that is changed is philosophical commentaries on mathematics, commentaries such as "Either the continuum hypothesis is determinately true or its negation is determinately true". (Field 1998a, p. 295)

Leeds objected to this view in the paper that I began with: his position, as one might guess, is that it is simply incoherent to accept set theory while regarding its concepts as indeterminate: especially (I think he would add) when one regards classical logic as fully applicable to them. And I have to agree that there can seem something fishy in the idea: if the philosophical commentaries about determinateness are as divorced from mathematical practice as the quotation may suggest, then they seem suspiciously idle. I think though that the position developed in this paper shows that the commentaries about determinateness needn't be divorced from mathematical practice. There can be a difference in mathematical practice between a set theorist who adopts a view in which the degrees of belief in the continuum hypothesis and its negation add to 1 and one who doesn't: the former may worry "Which is it?", and try to gather mathematical evidence in support of one answer or the other; the latter will think that talk of evidence is misguided, and will think that each possibility is worth developing. Of course, practical and aesthetic considerations may make some

answers more interesting than others; but these are not evidence. (For one thing, practical considerations may favor using a set theory that gives one value to the continuum in one context and using another in a different context. There is no problem in doing this as long as the contexts are kept distinct, but this makes clear that the utility of a mathematical theory in applications is quite independent of its “truth”. Similarly, distinct aesthetic criteria may point in different directions: if we keep truth out of the picture there is no need to ask which of two values of the continuum is prettier overall, anymore than there is a need to decide whether Bach or Beethoven is prettier overall.)

I should close by mentioning that Leeds himself offers a solution to Williamson’s puzzle (quite different from the one I have suggested). His discussion is brief, but the position seems to be that while some sense can be made of saying that it is indeterminate whether a given atom is part of Clinton’s body or whether a given person Joe is rich, what we are doing when we say that is (i) excluding the terms ‘Clinton’s body’ and ‘rich’ from the subpart of our language that (in that context) we take to be “first class”, and (ii) restricting the primary (disquotational) use of ‘refers’ and ‘true of’ to that subpart of our language. The indeterminacy of ‘rich’ and ‘Clinton’s body’ is then viewed as simply a matter of there being no uniquely best way to translate those terms into the subpart of our language to which we are applying the disquotation schema. I think this idea for trying to reduce all indeterminacy to indeterminacy of translation while recognizing vagueness in our ordinary language is ingenious, though I am skeptical of the possibility of ultimately making sense of it.²⁰

But even if it is workable, and preferable to the solution I have advanced, I doubt that it serves his ultimate aim, which involves excluding the possibility of there being any indeterminacy in our own terms that is more exciting than vagueness. Without further clarification in his proposal, I don’t think it is clear whether it excludes some or even all of the indeterminacy one might otherwise imagine in our own *scientific* terms. But in the case of the Brandom example I think it is clear that indeterminacy is not excluded: we can take our “first class language” to simply say that -1 has two square roots, without introducing a name like ‘ i ’ to “stand for one of them”.

Leeds actually does not discuss the Brandom example, so I am not sure that this conclusion would be unwelcome to him. But he certainly discusses the continuum example; and I think that something similar to what I just said for the Brandom example can be said there. More fully, the

idea is that we can in a sense accept set theory without taking the language of set theory as “first class”. We can do so even while being platonists in the sense of believing that there are non-physical eternally existing objects, and taking set theory to be talking about some of them. The idea would be to take as fundamental a theory T that doesn’t employ set-theoretic vocabulary but merely says that there are infinitely many non-physical eternally existing objects, and that postulates the consistency of basic set theory. (Consistency would have to be regarded as a primitive notion governed by its own axioms, rather than explained in set-theoretic terms; but there are independent reasons for adopting this attitude.) We then translate the language of set theory into T by taking ‘set’ to be true of some or all of the non-physical eternally existing objects and by interpreting ‘member of’ in any way that makes the usual axioms come out true. There are multiple ways of doing this, and different ones make different sentences about the size of the continuum come out true. So we get a multiplicity of translations between the language of set theory and the fundamental platonistic theory, and the continuum hypothesis comes out lacking in determinate truth value.²¹

One conclusion then is that even if we were to concede to Leeds his own solution to the Williamson problem, we cannot exclude the possibility of substantial indeterminacy in our own language, going far beyond mere vagueness. But I don’t see any reason to concede to Leeds his solution: I think that the model of vagueness and indeterminacy that I suggested in earlier sections is more attractive, and that it answers the skeptical doubts of both Leeds and Williamson.²²

Notes

1. Here and in what follows I ignore the fact that words like ‘rich’ are highly context-sensitive; I will imagine a context of use that settles many questions about who is rich. But any realistic such context will still leave many questions about who is rich unsettled.
2. A more moderate view would be that we can have an unrelativized notion of reference extending beyond our terms, it will just be highly indeterminate. But the more moderate view is not available to one who wants to deny indeterminacy in our own terms, for it postulates an indeterminacy in our term ‘refers’; that is why I have tentatively ascribed to Leeds the view that ‘refers’ be inapplicable to other terms except by relativization. Note that this less moderate view

requires that we have no notion of intertheoretic synonymy that could be used to break the relativization: intertheoretic synonymy is not merely indeterminate, it makes no sense.

(Indeterminacy of translation would be allowed, but it would have to be viewed as indeterminacy of *how to* translate, not as indeterminacy of “the synonymy relation”.)

3. A not implausible view for the semantics of the language of *natural* numbers is that it is completely indeterminate which ω -sequence of abstract objects our terms ‘0’, ‘1’, ‘2’, etc. stand for. If this view is accepted, it presumably extends to complex numbers, so that any abstract object whatever is a candidate for the reference of ‘i’ when conceived as part of a structure isomorphic to the complex numbers. However, there is a way to oppose the claim of indeterminacy in the natural number case: one can claim that each numeral determinately refers to a position in an ω -sequence. What is interesting about the example of the square roots of -1 is that this structuralist move won’t work: the position occupied by i is not structurally different from that occupied by -i, so even the move to structures won’t restore determinacy.

4. And of course we can avoid the indeterminate language here if we feel compelled to, by replacing talk of i and -i by sentences of the form “Whenever x is a root of -1 and y its conjugate, ...”.

5. Both ideas were articulated in Field 1974, but I missed the obvious point that they could be combined in a way that preserves schema (R). That point is spelled out in slightly more detail in Field 1998b.

6. The kind of details I have in mind are about what is going on sufficiently inside the boundaries of the black hole that any unfortunate individual who entered the black hole would have been destroyed by the tidal forces long before coming into a position to observe them; and they are details that could not be predicted from the prior state of the universe outside (or near the border of) the black hole, given the assumed indeterminism.

7. Something like this is arguably so for the radical theorist as well: although he has no need of a special determinately operator, he does use a different set of logical constants, and in the most plausible form of the view that I know of he must appeal to special contextual shifts in standards of acceptability; and it may be that an analogous problem arises in explaining some of these things.

I will not explore this.

8. If the language contains context-sensitive elements one must speak instead of probability assignments to sentences *relative to a context*; or probability assignments to propositions, where those are viewed as something like pairs of sentences and contexts. In what follows, I will assume a fixed context, so that we can speak simply of an assignment to sentences. When we come to vagueness, we won't be assigning *probabilities* to sentences, but we will be assigning other kinds of degrees of belief to sentences; again, this makes sense as long as the context is held fixed. (As remarked earlier, most vague terms are highly context-dependent, so the need to hold the context fixed is especially important.) If you like you can talk instead of assigning degrees of belief to propositions, as long as you don't employ a heavy duty notion of proposition on which 'Sam is rich' only expresses a proposition if Sam is not a borderline case.

9. Schiffer's theory, by contrast, has two distinct kinds of partial belief, operating according to different laws, one for where vagueness is allowed and one where it isn't; and the latter is not a special case of the former. For various reasons I think this unsatisfactory.

10. If we had started with a conditional probability function $P(X|Y)$ defined even when $P(Y|Z \vee \neg Z) = 0$, as in Popper 1959 (appendices iv and v), we could have defined Q by the rule $Q(X|Y) = P(DX|DY)$; this would allow $Q(X|Y)$ to be defined even when $Q(Y|Z \vee \neg Z) = 0$. The Popper approach, extended to the present context, requires that $P(X|Y)$ be 1 when Y is inconsistent with the modal system M , but allows other values when $P(Y|Z \vee \neg Z)$ is 0 without Y being inconsistent; so the same will be true of Q .

11. If one assumes that the probability of my trying out for the Yankees is strictly 0 rather than merely less than 10^{-10} , one needs the generalization mentioned in the previous footnote.

12. I should note that the conclusion that the Q -values of the Adams conditionals is 1 would not have been forthcoming had I started from a typical probability function based on a modal logic weaker than $S4$.

13. Indeed, we can do better still: any *one* of the occurrences of 'D' in the term on the left hand side can be dropped.

14. In particular, most of his theory concerns how belief functions are to be combined (“Dempster conditionalization”), and this seems to have no bearing on vagueness or indeterminacy.

15. $Q(DA \vee D\neg A) = P(D(DA \vee D\neg A)) = P(DA \vee D\neg A) = P(DA) + P(D\neg A) = Q(A) + Q(\neg A)$.

16. Let A be ‘This coin will come up heads on the next flip’, and C be any sentence that is intuitively independent of A . If the agent takes A and C as independent and takes the coin as fair, his set of probability functions on the sublanguage generated by A and C will be the set of P such that (i) $P(A) = 1/2$, (ii) $\acute{a} \leq P(C) \leq \hat{a}$ for some \acute{a} , \hat{a} (his “lower and upper probabilities”), and (iii) $P(A\&C) = P(A) \cdot P(C)$. Let B be $A \equiv C$. For each P in the set, we have that $P(B)$ is $1/2$, $P(A\&B) = 1/2P(C)$, and $P(A\vee B) = 1 - 1/2P(C)$; so $G_S(A)$ and $G_S(B)$ will each be $1/2$, $G_S(A\&B)$ will be $\acute{a}/2$, and $G_S(A\vee B)$ will be $1 - \hat{a}/2$, which violates (#) when \hat{a} is strictly greater than \acute{a} .

17. Suppose that A and C are intuitively independent sentences, but that this time the agent can put no lower or upper bounds (except 0 and 1) on the probability of either. That is, the agent adopts the set of probability functions (on the sublanguage generated by A and C) meeting condition (iii) of the previous note. In this case, for sentences E in the sublanguage, $G_S(E)$ is 1 if E is a tautology and 0 otherwise, just as it would have been if condition (iii) hadn’t been imposed: the shift from S to G_S throws away the crucial information that the agent takes A and C to be independent.

18. See Skyrms 1984, Ch 3.

19. If one thinks that the only possible indeterminacy here is in ‘set’, one may suppose that we can avoid postulating indeterminacy by taking the extension of ‘set’ to be “as large as possible”: see Leeds 1997, footnote 24, for such a suggestion. But in fact, an indeterminacy in the size of the continuum would be due at least as much to an indeterminacy in the membership predicate as to an indeterminacy in ‘set’. After all, the idea behind the argument for indeterminacy is that there is little to determine what the set-theoretic primitives are true of besides the axioms we accept, and ‘ ϵ ’ is one of those primitives. (I say ‘little’ rather than ‘nothing’, for reasons discussed in Field 1998a. But the other determinants don’t help in making the size of the continuum determinate.)

20. Perhaps the most central problem that the approach faces is that of distinguishing between indeterminate terms and terms that we would all be inclined to say are perfectly determinate but

simply differ in extension from terms in the “first class” part of our own language. (The problem is especially severe when the determinate terms have multiple “approximate translations” into our first class language, none clearly best, and these “approximate translations” differ in extension among themselves.) The obvious way for a disquotationalist to try to handle terms in other languages that differ in extension from terms in our language involves imagining adding those terms (or others with the same conceptual role) to our language, thereby allowing us to disquote. (See Quine 1953, p.135; Field 1994a, pp. 273-4.) But if we can imagine adding (analogs of) intuitively determinate foreign terms to our “first class language”, we can equally imagine adding (analogs of) terms like ‘/’ and ‘\’, or terms like ‘rich’. The proposed solution to the problem of saying what is defective about ‘/’ and ‘\’ and ‘rich’ threatens to disappear.

21. There is a complication: to handle the applications of mathematics to nonmathematical domains like physics and psychology, we need that these set theories are not merely consistent but jointly consistent with our consistent nonmathematical theories. And we must also suppose that each theory T of physics or psychology or whatever formulated in our ordinary set theory could be reformulated using the alternative platonistic theory. One way to do so would be to replace it by a nominalistic theory. Another would be to replace it by the claim that all of the nominalistic consequences of T-plus-some-chosen-set-theory are true. This second route may seem like a cheap trick. Note though that on the second route, the chosen set theory needn't be one that decides the size of the continuum, though it could if that turned out to be useful in the particular physical or psychological theory in question. Also, the set theory chosen for one physical or psychological theory needn't be compatible with the one chosen for another: this makes clear that the truth of the set theory is not being assumed in the superior conceptual framework, only its instrumental utility in a particular application.

22. I am grateful to Stephen Schiffer and Paul Horwich for a number of discussions, and to Schiffer and Tim Maudlin for helpful comments on a distant ancestor of this paper. Previous versions tentatively defended a different solution to “Williamson’s puzzle”. I am especially grateful to an editor and a reviewer at *Nous* for their skeptical response to that solution; their skepticism provoked the solution offered here. And David Barnett pointed out an error, and suggested several improvements, in the near-final draft.

