

“No Fact of the Matter”

Hartry Field

New York University

1. Introduction. There are many questions for which it is natural to think that there is no fact of the matter as to which of several competing answers is correct. Consider the question of when a human life begins: at birth? at conception? somewhere in between? This seems more a matter of legislation than of fact (and indeed the US Congress seems set to legislate it). Even if it were legislated, non-factuality would remain: if it were legislated that life begins at birth, well, births typically take a while; just what proportion of the fetus must have emerged? (And just where must any part of the fetus be to count as having emerged?) Similar questions would arise if it were legislated that life begins at conception (the likely legislation from the US Congress, it won't surprise you to learn): conception, like birth, is not an instantaneous affair and is a matter of degree.

So during which nanosecond did Jerry Falwell's life begin? The matter seems indeterminate, in the sense of, not a factual question. That seems to be the case both on our present conception of human life and on any future conception that is likely to result by (not necessarily Congressional) legislation.

This conclusion has been challenged by Timothy Williamson (1994); he thinks there is a fact of the matter, but unfortunately beings like us can never know the answer.¹ If there is an omniscient god, he knows, since by definition such a god would know all the facts. If there is a committee of omniscient gods, they all agree as to the nanosecond in which Falwell's life began. Even without introducing omniscience, there is no clear reason why more intelligent beings from other planets couldn't know the nanosecond in which Falwell's life began.

Here I don't mean that they could have a more precise concept of life than we have, call it life*, and could know during which nanosecond Falwell's life* began; that seems unproblematic. Rather, what I mean is that (on Williamson's account) these other beings could know the nanosecond in which his life began on *our* concept of life. And I don't mean that they could have a more precise concept of the referent of terms than we have, call it referent*, and could know the nanosecond that marked the beginning of the referent* of our term 'Falwell's life'. The times included in Falwell's life are the same as the times included in the referent of our term 'Falwell's life', on *our* concept of reference. That means that Williamson's view requires that extraterrestrials might know the first nanosecond in the referent of our term 'Falwell's life' on *our* account of reference, not on theirs.

Williamson's view, then, is that *we* can't know the nanosecond during which Falwell's life began, but other beings could know the answer to this question, not merely to some more precise substitute for this question. I'm not sure the grounds on which he can be confident that scientific advances couldn't lead to discoveries of such facts, if only the funding were available.

Williamson's view strikes many of us as beyond belief. Nonetheless, he has launched an impressive attack on those who would dissent from it. The initial argument is very simple. For any natural number N , let 'time N ' mean ' N nanoseconds after noon Eastern Standard Time on October 1 1933'. (I will pretend that phrase to be precise.) By the law of excluded middle (standardly called LEM, which is nice because it allows the characterization of intuitionism as Dutch LEM disease), we get each instance of the following schema:

(1P) (Falwell's life had begun by time N) \vee \neg (Falwell's life had begun by time N).

From a finite number of these plus the fact that Falwell's life hadn't begun by time 0 plus the fact that it had begun by time 10^{18} , plus the fact that for any N and M with $N < M$, if Falwell's life had begun by time N then it had begun by time M , a minimal amount of arithmetic and logic yields that

(F) There is a unique N_0 such that Falwell's life had begun by time N_0 and not by time $N_0 - 1$.

But then it seems that there is a fact of the matter as to which nanosecond his life began, viz. that between time $N_0 - 1$ and time N_0 (including the latter endpoint but not the former). That is the initial challenge.

One response to this initial challenge is to accept the derivation of (F), but to argue that despite (F) we can still make sense of the idea that the question of which nanosecond Falwell's life began is in some sense "nonfactual" (and not merely unknowable). Perhaps the best way to develop this response would be to concede that there is a "thin" or "pleonastic" sense of 'fact' in which if (F) holds then it is a fact that Falwell's life began during the nanosecond $(N_0 - 1, N_0]$; but to say that the more philosophically interesting notion is that of determinate fact, and the argument does not show that there is any determinate fact as to which nanosecond Falwell's life began, nor can it be generalized to show this. (The most famous version of this response is supervenualism, which I will discuss in a moment.) But there is a second stage of Williamson's challenge, which is to undermine this response.

The attempt to undermine it has two components. First, Williamson challenges the advocate of this response (who I'll call the classical indeterminacy theorist) to give an intelligible meaning of 'determinately' that makes it distinct from 'knowably'. The usual supervenualist version of the classical indeterminacy theory explains determinateness semantically, somewhat as follows: it is indeterminate whether Falwell's life began during nanosecond N because there are multiple legitimate precisifications of our term 'Falwell's life', and different of these legitimate precisifications have different nanoseconds as their beginnings. I take it that Williamson's response would be that this merely defines determinateness in terms of legitimate precisification, but that that notion

already has the idea of indeterminateness built in. A precisification of ‘Falwell’s life’ is presumably an entity located in a precise temporal interval (let’s say a closed interval, to avoid irrelevant complications); and what makes a precisification with initial point t_1 *legitimate* presumably includes that it is not definitely the case that Falwell’s life began prior to t_1 and not definitely the case that it began after t_1 . “Explaining” determinateness by means of the notion of legitimate precisification is no explanation at all. Once this is seen, I think it becomes highly plausible that there is no possibility of any reductive explanation of determinateness that captures the intent of the classical indeterminacy theorist. (I’m taking the explanation of indeterminacy as unknowability as violating the intent of the classical theorist.)

The best hope for the classical indeterminacy theorist, I think, is to concede that we can’t reductively explain the notion of determinateness in other terms, but to argue that we can nonetheless say a good bit to clarify the principles that govern the notion for those who already have it, and perhaps to make it intelligible to those who don’t yet have it.

But even if this can be done, there is a second component of Williamson’s response to the classical theorist. Namely, Williamson argues that introducing the notion of determinateness is simply beside the point. We began with the intuitive idea of there being no fact of the matter as to the nanosecond Falwell’s life began. We’ve now retreated to the claim that there is no *determinate* fact of the matter as to the nanosecond his life began, and insisted that ‘determinate’ doesn’t mean ‘knowable’. But whatever it means, we don’t seem to have captured enough of the idea of nonfactuality to be satisfying. For given that we’ve accepted the claim (F) that there is a unique N_0 such that Falwell’s life began in the nanosecond $(N_0 - 1, N_0]$, it seems uninteresting to add that this N_0 doesn’t have the additional property of being such that his life *determinately* began during that nanosecond. Why should we care about the determinately property? And why should the fact that there is no N such that his life *determinately* began during nanosecond $(N - 1, N]$ keep us from wondering which N_0 is the value marking the nanosecond during which his life began? Why should it keep us from imagining aliens who know the answer to that question? Why should it keep us from being very worried about the possibility that some action I may have performed was done after the critical time N_0 rather than before? And so forth. Without answers to such questions, we have done nothing to capture the intuitive idea of non-factuality. And satisfactory answers to them are not easy to come by.

These reflections pose a serious difficulty for any attempt to resolve the problem within classical logic, and I have come to think that the difficulty is ultimately insurmountable.² If we are to resist Williamson’s conclusion of factuality everywhere, it seems that we must somehow block the derivation of (F).

Obviously one way to do this is to reject some instances of excluded middle. But if ‘reject’ means ‘deny’ (in the sense of, ‘accept the negation of’), there is a serious problem with this, as Williamson points out. Suppose we accept

(1N) $\neg[(\text{Falwell's life had begun by time } N) \vee \neg(\text{Falwell's life had begun by time } N)]$.

The expression in brackets is a disjunction, and surely on any reasonable logic a disjunction is weaker than either of its disjuncts. So denying the disjunction has got to entail denying each disjunct, and so accepting (1N) clearly commits us to accepting both of the following:

(2a) $\neg(\text{Falwell's life had begun by time } N)$

(2b) $\neg\neg(\text{Falwell's life had begun by time } N)$

But (2b) contradicts (2a), so (1N) has led to contradiction.

One response to this is Graham Priest’s “dialetheism” (1998). His formulation of the response is that we should accept some contradictions, adopting a “paraconsistent logic” on which contradictions don’t entail everything and so aren’t so bad. Many people are repelled by talk of accepting contradictions, but the substance of dialetheism can be put differently, as the claim that $\neg A$ doesn’t really contradict A . (In this alternative formulation, B would be said to contradict A only when their conjunction entails everything.)

Whichever way we prefer to put it, I do not think we ought to reject dialetheism out of hand. But there are problems with using it in response to Williamson. For one thing, since the dialetheist accepts (2a), and (1P) is a disjunction with (2a) as one disjunct, the dialetheist will accept (1P) as well as (1N): (1P) follows from (2a) on any reasonable logic, including all the standard paraconsistent logics. But then we can argue from (1P) to Williamson’s conclusion that there is a unique nanosecond in which Falwell’s life began, in precisely the same way as before, so the conclusion has not been blocked. The conclusion has been *denied*—from (1N) we can conclude that there is not a unique nanosecond during which his life began³—but it has also been asserted. This “inconsistency” is not in itself a problem, it is just a further instance of dialetheist doctrine; but it is disappointing that we are left in a position of thinking that it is just as correct to assert that there is a fact of the matter as to the nanosecond in which Falwell was born as to assert that there is no fact of the matter.

In short, despite its denial of the relevant instances of excluded middle, the dialetheist approach fails because it does *not* block the derivation of (F)—it merely allows an additional derivation of its negation—and this is inevitable on any remotely reasonable logic that *denies* instances of excluded middle, i.e. accepts their negations. To block the derivation of (F), we must reject the relevant instances of excluded middle *without denying them*.

I should mention that there is an independent reason for thinking that a logic without excluded middle is the most natural logic to use when one wants to take into account the possibility of factual indeterminacy: the semantic paradoxes seem to point to a kind of factual indeterminacy in certain sentences involving ‘true’ and related terms, and there are strong reasons to think that a satisfactory solution of them will involve restricting excluded middle. (See Field 2003b for such a solution, Field 2003d for a comparison of it to classical solutions, and Field 2003c for a comparison to dialetheist solutions.) I’ll say a little bit about this below.

But there are still many questions to be answered. For instance,

#1. What exactly is it to reject a claim, if not to deny it? As we’ll see, rejection has to be stronger than mere non-acceptance; what is the additional element? This question does not seem to have been addressed by those who would “reject” some instances of excluded middle, but it very much needs an answer: e.g. we need to say what our difference in attitude is between those instances of excluded middle we accept, those we reject, and those we neither accept nor reject.

#2. If we give up excluded middle as a general law, what is the appropriate logic to replace it? (If the logic is to be appropriate to the semantic paradoxes, many popular proposals, such as “fuzzy logic” (based on Lukasiewicz continuum-valued semantics) and intuitionist logic, are ruled out.)

And most crucially,

#3. How exactly does the proposal help with the problem with which we began?

That problem, as I originally put it, was that we wanted to make sense of the claim that there is no fact of the matter about certain questions, such as whether Falwell’s life had begun by nanosecond N. But if ‘it is a fact that p’ is just a pleonastic equivalent of ‘p’, then the *rejection* of excluded middle does not allow us to do this: to assert that there is no fact of the matter as to whether his life had begun by nanosecond N would be equivalent to asserting (1N), which as we’ve seen leads to the *acceptance* of a fact of the matter (as well as acceptance of its negation). On the pleonastic interpretation of ‘fact’, we cannot literally assert that there is no fact of the matter: we can only reject there being a fact of the matter (in the sense of rejection to be explained).

If the best we could do were to allow for the rejection of there being a fact of the matter, then the approach I’m suggesting could be only a limited success. It would be *something* of a success: it is an important advance over the classical approach that we can explain the difference between the attitude most of us have toward statements like

(3) The largest whole number of nanoseconds between the start of Falwell’s life and the start of Ashcroft’s is even

and the attitude we have toward statements that seem perfectly factual even though their answers may be unknowable, such as perhaps

(4) Attila's maternal grandmother weighed less than 125 pounds on the day she died.

But there are those with different attitudes than ours toward these sentences—many right-to-lifers would accept the factuality of (3), and many verificationists would reject the factuality of (4)—and we need to be able to debate our attitudes; and without some means of “denying factuality” our ability to conduct such debates would be severely limited.

To resolve this problem I propose to borrow a solution from the classical theorist: introduce a determinately operator **D**. **DA** means ‘it is determinately the case that A’; so if we define **GA** as **DA** \vee **D** \neg A, **GA** means ‘it is determinate whether A’. Much more needs to be said about this operator, but the general idea is that though we cannot accept $\neg(A \vee \neg A)$, we can accept $\neg(\mathbf{DA} \vee \mathbf{D}\neg A)$, i.e. $\neg\mathbf{GA}$. Moreover, the acceptance of $\neg\mathbf{GA}$ should in some sense be a “projection” of our rejection of $A \vee \neg A$; of course, more will need to be said to explain this. Once we have such an operator, we can either interpret ‘it is a fact that A’ as equivalent to **DA** rather than as equivalent to A, or we can regard the intuition that “there is no fact of the matter” about questions like (3) as better expressed by the claim that there is no *determinate* fact of the matter.

If the non-classical theorist needs to follow the classical theorist in invoking a determinately operator, wherein lies the advantage of the non-classical approach? There are several advantages. First, the non-classical theorist can do a better job of explaining the operator. Second, by linking the acceptance and rejection of **DA** \vee **D** \neg A with the acceptance and rejection of $A \vee \neg A$, the nonclassical theorist can make more plausible that the operator being introduced has something to do with the intuitive idea of factuality. Third and most crucial, any relevance that the classical theorist may claim for the acceptance of $\neg(\mathbf{DA} \vee \mathbf{D}\neg A)$ to the intuitive idea of nonfactuality seems undercut by the classical theorist's acceptance of $A \vee \neg A$: e.g. if either Falwell's life began in an even numbered nanosecond or it didn't, why couldn't aliens know the answer, and why couldn't we wonder what the answer is? No such puzzle arises for the nonclassical theorist.

It is evident that to fill out the answer to #3 just sketched, much more will be need to said about the determinately operator and to the links between on the one hand accepting or denying **DA** \vee **D** \neg A and on the other hand accepting or rejecting $A \vee \neg A$. It is largely these questions, plus questions #1 and to a lesser extent #2, that will be the concern of the rest of the paper. Questions about higher order indeterminacy will need to be addressed along the way: it will turn out that several different phenomena go under this heading, and as a result the discussion of it is divided among the following three sections.

2. Rejecting instances of excluded middle. While the treatment of “non-factual discourse” that I will be advocating does use the notion of determinateness, the notion of determinateness appears rather late: the main moves are made before the notion of determinateness is introduced.⁴ In particular, it will be important to my purposes that the logical practices of rejecting certain instances of excluded middle while accepting others can be explained without using the notion of determinateness. Among the instances of excluded middle we presumably want to “reject” are certain instances of the schema (1P) (Falwell’s life had begun by time N) \vee \neg (Falwell’s life had begun by time N).

This presumably requires “rejecting” the disjuncts

(5a) Falwell’s life had begun by time N

and

(5b) \neg (Falwell’s life had begun by time N)

as well. But what is it to “reject” these claims? A defect in most treatments of indeterminacy that reject excluded middle is that they make no serious attempt to answer this.

We’ve already ruled out one answer: that rejection is acceptance of the negation. It isn’t just that accepting a counter-instance to excluded middle would require dialetheism. The more fundamental point is that rejecting a claim should preclude accepting it; since the point of dialetheism is that accepting $\neg A$ does *not* preclude accepting A , it is clear that even a dialetheist must distinguish between rejection and acceptance of the negation. So *anyone* who wants to reject instances of excluded middle, whether dialetheist or not, needs an account of rejection distinct from acceptance of the negation. A related point is that as noted before, there is a compelling reason for anyone who accepts the negation of (1P) to accept (1P) itself, and therefore *not* reject (1P). (From the negation of (1P) we infer the negation of its first disjunct; but that is the second disjunct of (1P), so we can infer (1P) itself.) So it isn’t that the view that we should deny a given instance of (1P) is *a bad version of* the view that we should reject that instance; rather, it simply isn’t a version, good or bad, of the view that we should reject that instance.

One might say that to reject A is to regard it as either false or indeterminate. But at this stage we have no explanation whatever of the notion of indeterminacy, so this seems unsatisfactory. An alternative is to say that to reject A is to regard it as either false or meaningless. But this is obviously inadequate to the case at hand: no instance of (5a) or (5b) is meaningless in any ordinary sense. (Even those instances with values of N corresponding to the actual “borderline region” are ones that someone might well firmly believe or firmly disbelieve, say if she had a false belief as to Falwell’s birthday; indeed we use the meaning of (5a) or (5b), together with empirical information, to decide

whether to believe (5a) or (5b) and whether to regard them as indeterminate.) Reichenbach 1938 calls claims with the general flavor of (5a) and (5b) “physically meaningless”. (To avoid presupposing a commitment to physicalism, we might modify this to “*metaphysically* meaningless”.) This seems fine, as long as one recognizes that “(meta)physical meaningfulness” is not meaningfulness in any ordinary sense; but it isn’t very helpful, in that it seems to be just another word for ‘indeterminate’.⁵ The same complaint can be raised against the claim that (5a) and (5b) “don’t express propositions”: if this is construed as compatible with meaningfulness (as it needs to be to be relevant here) then it seems like just a synonym for indeterminacy.⁶ To repeat, I don’t say it is *wrong* to equate rejecting A with believing it false or indeterminate, it is simply unhelpful. A better strategy is to give an independent explanation of rejection, and use that to help clarify the notion of indeterminacy. That will be my procedure.

Here’s another common proposal: to reject A is to regard A as *not true*; where ‘not true’ is supposed to be broader than ‘false’ (i.e. ‘has a true negation’), it is supposed to include things that are neither true nor false. But on the clearest notion of truth, the claim that A is true is simply equivalent to the claim A, so the claim that A is not true is equivalent to the claim $\neg A$; since we are rejecting not only the claims (1P), (5a) and (5b) but also their negations, we must reject the claim that (1P), (5a) and (5b) aren’t true; so we don’t accept those claims, so rejection can’t be acceptance of non-truth. Put another way: given the clearest notions of truth and falsity, we cannot (in the non-paraconsistent logics under consideration) say of any sentence that it is neither true nor false: the claim that A is false is equivalent to the claim that $\neg A$ is true, hence is equivalent to $\neg A$ itself, so asserting that A is neither true nor false is equivalent to asserting $\neg(A \wedge \neg A)$, which as we’ve seen is not allowed. More generally, we can only assert that a sentence is either false or lacking in truth value if we can assert that it is false. If to reject A were to believe it to be false or lacking in truth value, we couldn’t reject any instances of (1P), (5a) or (5b).

The previous paragraph assumes that $\text{True}(\langle A \rangle)$ is completely equivalent to A: that they are intersubstitutable in all contexts.⁷ I don’t doubt that one can introduce a notion for which this intersubstitutivity fails, and call it truth if one likes; indeed, given the above notion of truth (call it “weak truth”) and the notion of determinacy, one can introduce the notion of determinate weak truth (call this “strong truth”). If ‘True’ is used for strong truth instead of for weak truth, ‘ $\text{True}(\langle A \rangle)$ ’ will not be intersubstitutable with A in all contexts: the former will imply the latter but not conversely. So rejecting A might be identified with regarding A as not strongly true. However, this would be just another version of a previous proposal: rejecting A would be identified with regarding it indeterminate. For reasons already discussed, this is not very helpful.

One last obvious but inadequate account of rejection is: refusal to accept. The problem with this is that it is too weak to capture the intuitive idea that rejecting A should

roughly coincide with regarding it as false or indeterminate. My attitude toward ‘Falwell’s life began in an even-numbered nanosecond’ and its negation is quite different from my attitude toward ‘Attila’s maternal grandmother weighed less than 125 pounds on the day she died’ and its negation: I wouldn’t accept the latter sentence or its negation, for lack of evidence (and indeed think it extraordinarily unlikely that it is possible to gather enough evidence on this matter to make me accept either one), but I have no temptation to assert that the matter is indeterminate. (It *could* be indeterminate, e.g. if she was borderline dead at the midnight of a day on which she lost a critical pound; but I have no reason to suppose that it is.) The sense in which one who opposes LEM might reject (5a) and (5b) (for certain values of N) is stronger than mere refusal to accept them. The same goes for the instances of (1P). Consider an opponent of LEM who has no idea how old Falwell is. Such a person may *fail to accept* those instances of (1P) where N corresponds to times around his birth or his conception or somewhere in between; but does she *reject* such instances? In the sense of rejection that is relevant to the notion of indeterminacy and will ultimately help ground it, she neither accepts nor rejects them: she has no idea whether they correspond to determinate cases.

So we need another account of rejection. The key to providing one is to recognize that the refusal to accept all instances of excluded middle forces a revision in our other epistemic attitudes. A standard idealization of the epistemic attitudes of an adherent of classical logic is the Bayesian one, which (in its crudest form at least) involves attributing to each rational agent a degree of belief function that obeys the laws of classical probability; these laws entail that theorems of classical logic get degree of belief 1. Obviously this is inappropriate if rational agents needn’t accept all instances of excluded middle. But allowing degrees of belief less than 1 to some instances of excluded middle forces other violations of classical probability theory.

In particular, consider the following three laws of classical probability theory:

$$(6) \quad P(A \vee B) + P(A \wedge B) = P(A) + P(B);$$

$$(7) \quad P(A \wedge \neg A) = 0;$$

$$(8) \quad P(A) + P(\neg A) = 1.$$

Instantiating B with $\neg A$ in (6), these clearly imply that $P(A \vee \neg A) = 1$. The opponent of LEM can’t accept that for degrees of belief (at least, not as long as degrees of belief are confined to the interval $[0,1]$, as I shall assume), and so he must give up one of (6)-(8). If such an opponent renounces dialetheism (the assertion of $A \wedge \neg A$), he will presumably adhere to (7), and has no obvious reason to give up (6). (8), on the other hand, seems quite out of keeping with the view;⁸ we need to replace the ‘=’ in it by ‘ \leq ’ (and add the law that $P(\neg\neg A) = P(A)$).

The relevance of this to acceptance and rejection is that accepting A seems intimately related to having a high degree of belief in it; say, a degree of belief at or over a certain threshold $T > 1/2$.⁹ So let us think of rejection as the dual notion: it is related in the

same way to having a low degree of belief, one at or lower than the co-threshold $1-T$. In the context of classical probability theory where (8) is assumed, this just amounts to acceptance of the negation. But without (8), rejection in this sense is weaker than acceptance of the negation (while remaining stronger than failure to accept, in that sentences believed to degrees between $1-T$ and T will be neither accepted nor rejected).

The natural way to employ this in the context of indeterminacy is to suppose that for a person who we are inclined to describe as “certain of the indeterminacy of A ”, $P(A)$ and $P(\neg A)$ will both be 0 (where P is the person’s degree of belief function); equivalently (since degrees of belief are never negative), $P(A) + P(\neg A)$ will be 0. And a person “certain of the determinacy of A ” will be one for whom $P(A) + P(\neg A) = 1$. More generally, it is natural to take $P(A) + P(\neg A)$ —or equivalently, $P(A \vee \neg A)$ —as a measure of the extent to which the agent “believes A determinate”. Belief revision on empirical evidence goes just as on the classical theory (by conditionalizing); this allows the “degree of certainty of the determinacy of A ” to go up or down with evidence (as long as it isn’t 1 or 0 to start with).

(What I’m calling “degree of belief” might be called “lower degree of belief”, and the function assigning it denoted P_* ; one could regard $1-P_*(\neg A)$ as the “upper degree of belief” in A , and call it $P^*(A)$. $P_*(A)$ is always less than or equal to $P^*(A)$; equality obtains precisely when $P_*(A)+P_*(\neg A)=1$, i.e. when the agent is certain of the determinacy of A . Whereas $P_*(A)+P_*(\neg A)$ can never exceed 1, $P^*(A)+P^*(\neg A)$ can never be less than 1. Why have I used the term ‘degree of belief’ for $P_*(A)$ rather than $P^*(A)$, or rather than for the interval from $P_*(A)$ to $P^*(A)$? Given a non-paraconsistent logic where one can’t consistently accept both A and $\neg A$ at the same time, it would not be natural to use the unqualified term ‘degree of belief’ for upper degrees. But I concede that the intervals might be a more natural choice for “degrees of belief” than their lower bounds: certainty of the determinacy of A would then be the same as having a point-valued degree of belief in A . The matter however is purely terminological, and what I say here can easily be recast into that alternative terminology.)

So far I have for simplicity assumed the standard idealization, that each agent has a unique degree of belief function. This is obviously a considerable idealization: there is surely no determinate fact of the matter as to whether the fifth digit in the decimal expansion of my degree of belief that it will rain tomorrow is even. How do we handle this? Not, I think, by using more complicated sorts of mathematical entities than real numbers to represent our degrees of belief: e.g. sets of reals, or sets of sets of reals, or whatever. Doing that doesn’t really help, and in some ways makes the problem worse. Rather, the appropriate treatment of indeterminacy in our degrees of belief is the same as for indeterminacy elsewhere: we give up excluded middle. That is, the problem about whether the fifth digit in my degree of belief that it will rain tomorrow is even arises from assuming that it either is even or it isn’t. We should give that up. I don’t mean that we

need to develop a special mathematics of “numbers that can’t be assumed to be either even or not even”; the mathematics of degrees of belief can be assumed to be ordinary functions into $[0,1]$ satisfying (6) and (7) and the replacements for (8). What can be indeterminate is the relation between a believer and such a function, and so it is only when it comes to attributing such a function to an agent that excluded middle can fail. So we can’t assume excluded middle for certain claims about X ’s degrees of belief, even though the mathematics of the degree of belief functions themselves is perfectly classical. (Having made this point, I will sometimes ignore it in what follows for reasons of simplicity, though I will occasionally bring it in when it seems relevant.)

I remarked before that the semantic paradoxes seem to point to a kind of factual indeterminacy in certain sentences involving ‘true’ and related terms. If this is right, it seems to support the proposal that sentences thought of as indeterminate are such that both they and their negations are believed to degree 0. Consider the Liar sentence L , which asserts of itself that it is not true. Since L is equivalent to $\neg\text{True}\langle L \rangle$ which in turn entails $\neg L$, then L entails $\neg L$; L also entails itself, so L entails the contradiction $L \wedge \neg L$. Analogously, $\neg L$ also entails $L \wedge \neg L$.¹⁰ But if L entails something and $\neg L$ does too, then presumably $L \vee \neg L$ entails it (disjunction elimination); so $L \vee \neg L$ entails a contradiction. Now, what degree of belief should we have in $L \vee \neg L$? It would seem as if any positive degree of belief we had in it should extend to the contradiction which it entails. But (barring dialetheism) we should not believe a contradiction to any positive degree; so we must believe $L \vee \neg L$ to degree 0. (Similarly, we must believe each of the disjuncts to degree 0.) I regard the case of the Liar sentence as a good model for sentences that we are certain are indeterminate.

Returning to the general proposal about degrees of belief, it is that the degree to which an agent X regards A as determinate is measured by $P_X(A) + P_X(\neg A)$. (Because it may be indeterminate what X ’s degrees of belief are, this allows for a kind of indeterminacy in the extent to which X regards A determinate.) So far, the proposal takes disagreement about the determinateness of A as simply disagreement in attitude: disagreement about what degrees of belief to have. But it seems natural to give an alternative construal that takes “disagreement about the determinateness of A ” at face value: as disagreement about some proposition $\mathbf{G}A$. This requires the introduction of a determinacy operator \mathbf{G} into the language, and into the range of each agent’s degree of belief function; and it requires that $P_X(\mathbf{G}A) = P_X(A) + P_X(\neg A)$, so that the literal sense of “disagreement about the determinacy of A ” will pretty much coincide with the “disagreement in attitude” sense previously motivated. Actually, instead of taking \mathbf{G} as primitive, it’s simpler to take as primitive an operator \mathbf{D} , where $\mathbf{D}A$ means that it is determinately the case that A . The claim $\mathbf{G}A$ that A is determinate (i.e., that it is determinate *whether* A) is the claim that $\mathbf{D}A \vee \mathbf{D}\neg A$.¹¹

The details of the probabilistic laws governing **D** and hence **G** will depend on the details of the **D**-free logic, a matter to be deferred until the next section. But as already remarked, for **G** to really be an operator meaning determinacy, it seems that we need to have acceptance of **GA** coincide with acceptance of $A \vee \neg A$, and analogously for rejection. Since **GA** is just $\mathbf{DA} \vee \mathbf{D}\neg A$, it seems that the only reasonable way to achieve this is to demand that $P(\mathbf{DA})$ must equal $P(A)$, where again P attributes “lower” degrees of belief.¹² The extra strength that **DA** has over A will come out in $P(\neg\mathbf{DA})$ sometimes being greater than $P(\neg A)$ (i.e. in the *upper* degree of belief in **DA** sometimes being lower than that in A). More generally, $P(\neg\mathbf{DA} \wedge \neg\mathbf{DB})$ can be greater than $P(\neg A \wedge \neg B)$; since $\neg\mathbf{GA}$ is equivalent to $\neg\mathbf{DA} \wedge \neg\mathbf{D}\neg A$, this is what allows us sometimes to assert $\neg\mathbf{GA}$ even though we can never assert $\neg A \wedge \neg\neg A$.

I’ve proposed an identity between X ’s degree of belief in **GA** and his degree of belief in $A \vee \neg A$. (Correspondingly, between his degree of belief in **DA** and his degree of belief in A .) To what extent is there a link between X ’s degree of belief in $\neg\mathbf{GA}$ and his degree of belief in $A \vee \neg A$? Obviously there is at least this much of a link: $P_X(\neg\mathbf{GA})$ can’t be more than $1 - P_X(A \vee \neg A)$. But it is hard to see what more than this we can say, given the possibility of higher order indeterminacy. (That is the explanation of the hedge phrase “pretty much” two paragraphs back.)

To elaborate: there seem to be at least two different ways in which one might regard an agent’s degrees of belief as showing a commitment to something reasonably called higher order indeterminacy (and I’ll note a third later on). The first (HOI_1) arises prior to the introduction of the determinateness operator, and has already been mentioned: since an agent’s degrees of belief are themselves indeterminate, it may be indeterminate how high $P_X(A \vee \neg A)$ is;¹³ hence (once we do introduce the determinateness operator), it may be indeterminate how high $P_X(\mathbf{GA})$ is. The second way (HOI_2), which seems formally independent of the first, is that the agent might not adhere to the instance of excluded middle $\mathbf{GA} \vee \neg\mathbf{GA}$: $P_X(\mathbf{GA} \vee \neg\mathbf{GA})$ (insofar as it is determinate what this is) might be low, and hence $P_X(\neg\mathbf{GA})$ might be lower than $1 - P_X(\mathbf{GA})$, i.e. than $1 - [P_X(A) + P_X(\neg A)]$. In that case, an agent could treat A as indeterminate in the original sense (low $P_X(A) + P_X(\neg A)$) without this being fully reflected in high $P_X(\neg\mathbf{GA})$, but only in low $P_X(\mathbf{GA})$. Assertion of $\neg\mathbf{GA}$ isn’t proper unless $P_X(A) + P_X(\neg A)$ is low; but if there is enough higher order indeterminacy of this second sort, $P_X(A) + P_X(\neg A)$ could be low and yet it be improper to assert $\neg\mathbf{GA}$, though improper to assert **GA** as well. (Here I am indebted to Richard Dietz, who noted that I did not consistently take account of the point in Field 2000.)

It is tempting to try to limit the scope of higher order indeterminacy of the second kind, by linking it somehow to higher order indeterminacy of the first kind. That is, it is tempting to suppose that one should attach a low degree of belief to $\mathbf{GA} \vee \neg\mathbf{GA}$ only to

the extent that it is indeterminate what one's degree of belief in $A \vee \neg A$ is. There is a bit of an issue as to how this idea would be best formalized; and I'm not in the end convinced that it is a good idea anyway. I will not pursue it. It does seem that a good theory ought to have *something* to say about sufficient conditions for asserting $\neg GA$. In Section 4 I will consider the possibility of defining **D** and hence **G** in terms of a more primitive operator about which more can be said, and one advantage of so doing is that it would give stronger constraints on when the assertion of $\neg GA$ is allowed.

It should be clear from this last discussion that we cannot completely determine the probability function for the language that includes **D** from the probability function for the **D**-free language that is built up from atomic predicates using only the Kleene connectives. That isn't all that surprising: we can't determine the probability function for the language with \wedge and \vee from the probability function for the fragment that excludes these; why should we expect such determination for **D** when we don't have it for \wedge and \vee ? Even without such complete determination, the probabilistic laws governing **D**, including (but not limited to) those that relate **DA** to **D**-free formulas, do a great deal to clarify the meaning of **D**: in particular, the fact that $P(\mathbf{DA})=P(A)$, and that $P(\mathbf{DA} \vee \mathbf{D}\neg A) = P(A \vee \neg A)$ (and hence equals $P(\mathbf{D}(A \vee \neg A))$), seem very important to understanding the notion.¹⁴

3. Kleene logic and semantic value. I've said little so far about the logic I'm advocating, other than that it not include excluded middle but not include any negations of instances of excluded middle either. But I will now assume, without argument, that the logic of the connectives *other than the conditional* should be the logic K_3 of the strong Kleene 3-valued truth-tables, taking the valid inferences to be those that preserve the top value in every valuation.¹⁵ (So a good degree of belief function should *validate* K_3 , in the sense that if A_1, \dots, A_n entail B in K_3 then $P(A_1 \wedge \dots \wedge A_n) \leq P(B)$.)

More fully, we think of sentences as having a "semantic value" of 1, $\frac{1}{2}$ or 0. (For the moment it will do no harm to assume excluded middle for the assignment of semantic values to sentences, though I will later question this assumption.) Given the values of atomic sentences, the values of complex sentences built up using \neg , \wedge and \vee is given as follows:

A	$\neg A$
1	0
$\frac{1}{2}$	$\frac{1}{2}$
0	1

$A \wedge B$	B=1	B= $\frac{1}{2}$	B= 0
A=1	1	$\frac{1}{2}$	0
A= $\frac{1}{2}$	$\frac{1}{2}$	$\frac{1}{2}$	0
A=0	0	0	0

$A \vee B$	B=1	B= $\frac{1}{2}$	B= 0
A=1	1	1	1
A= $\frac{1}{2}$	1	$\frac{1}{2}$	$\frac{1}{2}$
A=0	1	$\frac{1}{2}$	0

The values of quantified sentences is determined analogously.¹⁶

How are we to interpret the values 1, $\frac{1}{2}$ and 0 that figure in these tables? Primarily, I'd say, by their role in an account of commitment and validity. Committing oneself to a claim is roughly the same as thinking it has value 1 (and believing a claim to a certain degree is roughly the same as believing to that degree that it has value 1). Note that I've characterized validity in a relatively weak sense: preservation of the value 1. On this notion of validity, an inference is valid when commitment to the premises brings implicit commitment to the conclusion; but it is not required that commitment to the negation of the conclusion brings commitment to negation of the conjunction of the premises. I think that the notion of validity defined by this stipulation together with the rules for the connectives is a very natural one. I'll mention just one of its consequences, the rule of disjunction elimination:

(\vee E) If the inferences from A to C and from B to C are both valid, then so is the inference from $A \vee B$ to C.

For if the inference from $A \vee B$ to C were invalid, there would be a valuation in which $A \vee B$ had value 1 and C didn't; but then either A would have value 1 and C not, or B would have value 1 and C not, so either the inference from A to C or the inference from B to C would be invalid. I take it that our reasoning practices do involve the acceptance of (\vee E); the Kleene tables (together with the definition of validity) provide a nice characterization of these reasoning practices. I think that the consequence of the Kleene tables for our deductive practice is enough to "give meaning to" the values 1, $\frac{1}{2}$, and 0.

But some may be tempted to translate the colorless labels 1, $\frac{1}{2}$ and 0 into more familiar language. One way *not* to translate them is to say that 1 means "true", 0 means "false" (i.e. "has a true negation"), and $\frac{1}{2}$ means "neither true nor false". For we presumably want to allow asserting of a given sentence (say about which nanosecond

Falwell's life began) that it has value $\frac{1}{2}$. But on our preferred use of 'True' (the use with no notion of determinacy built in), $\text{True}\langle A \rangle$ is intersubstitutable with A ; this makes ' $\neg\text{True}\langle A \rangle \wedge \neg\text{True}\langle \neg A \rangle$ ' equivalent to ' $\neg A \wedge \neg\neg A$ ', which is never assertable (except for the dialetheist, whose views I've already put aside). ' $\neg\text{True}\langle A \rangle \wedge \neg\text{True}\langle \neg A \rangle$ ' is the same as ' $\neg\text{True}\langle A \rangle \wedge \neg\text{False}\langle A \rangle$ '; so we can never assert of sentences that they are not true and not false, so we shouldn't interpret attributions of value $\frac{1}{2}$ as having that meaning.

A somewhat better way of understanding the values 1, $\frac{1}{2}$ and 0 is that 1 means "determinately true", 0 means "determinately false" (i.e. "has a determinately true negation"), and $\frac{1}{2}$ means "neither determinately true nor determinately false". But this may not be quite accurate: it would fail if we both posit a third kind of higher order indeterminacy, beyond the two already mentioned, and also adopt a certain convention for how to talk about it. This will require a few paragraphs of explanation.

I take it that just as there are times t for which it is indeterminate whether Falwell's life began before t , so too there are times t for which it is indeterminate whether Falwell's life *determinately* began before t : our concept of determinateness has indeterminacies in it, just as does our concept of a life. I've argued that regarding it indeterminate whether Falwell's life began before t involves rejecting the following instance of excluded middle:

(Falwell's life began before t) \vee \neg (Falwell's life began before t).

Similarly, regarding it indeterminate whether Falwell's life *determinately* began before t would seem to involve rejecting the following instance of excluded middle:

(Falwell's life determinately began before t) \vee \neg (Falwell's life determinately began before t).

This is the second of the two construals of higher order indeterminacy already considered: we have higher order indeterminacy in this sense (HOI_2) when it is inappropriate to apply excluded middle to claims about determinacy. My own view is that any notion of determinacy worthy of the name will admit HOI_2 : excluded middle will sometimes fail for attributions of determinacy. (Not only is this plausible for ordinary sorts of indeterminacy, it seems especially compelling for anyone who wants to treat the semantic paradoxes as due to indeterminacy. Suppose that instead of the ordinary Liar sentence L , which says of itself that it is not true, we consider the modified Liar sentence L^* , which says of itself that it is not determinately true. Then $\neg L^*$ is equivalent to \mathbf{DL}^* , which surely implies L^* ; so $\neg L^*$ implies $L^* \wedge \neg L^*$, a contradiction. But it is natural to suppose that for any sentence A , A implies \mathbf{DA} (in the weak sense of implication which simply registers transmission of commitment from premise to conclusion, not from negation of conclusion to negation of premise). If so, then L^* implies \mathbf{DL}^* , which is equivalent to $\neg L^*$; so L^* also implies the contradiction $L^* \wedge \neg L^*$. By disjunction elimination, $L^* \vee$

$\neg L^*$ implies this same contradiction, so it had better not be a logical truth. But it is equivalent to $DL^* \vee \neg DL^*$, so that had better not be a logical truth either.)

Now, accepting HOI_2 does not undermine taking semantic value 1 to mean “determinately true” (and 0 to mean “determinately false” and $\frac{1}{2}$ to mean “neither determinately true nor determinately false”): the only import is that *if we do so read these notions, we must be careful not to assume excluded middle for attributions of semantic value*. We should not assume, for instance, that for every sentence A, it either has semantic value 1 or it doesn’t. Indeed, assuming excluded middle for attributions of semantic value would be unreasonable, if semantic value is to be anything but a highly artificial technical notion. Here are two arguments for this claim:

(1) Commitment to excluded middle for assertions about semantic value would imply, by slight adaptation the Williamson argument, that there is a unique N_0 such that ‘Falwell’s life had begun by time N_0 ’ has semantic value 1 while ‘Falwell’s life had begun by time $N_0 - 1$ ’ doesn’t; unless ‘has semantic value 1’ is taken to be a highly artificial technical notion, this seems absurd.

(2) Consider the sentence that asserts of itself that it does not have semantic value 1; the assumption that it has value 1 seems to lead to contradiction, as does the assumption that it doesn’t, so the most obvious way to block the threatened paradox is to refuse to grant that it either has value 1 or doesn’t. (I don’t deny that for many purposes—in particular, for set-theoretic consistency proofs—it may be important to introduce artificial notions of semantic value that do obey excluded middle. On these artificial notions, the paradox must be blocked in some other way. For discussion, see the last section of Field 2003b or Section 8 of Field 2003d.)

However, there is some temptation to accept higher order indeterminacy in a more controversial sense (HOI_3): roughly, a sense in which something might fail to be *determinately* determinately the case without failing to be (merely) determinately the case; more accurately, a sense in which $\neg DDA$ doesn’t imply $\neg DA$. If we do accept higher order determinacy in that more contentious sense (and I don’t take it to be obvious that we should), there’s a question as to whether we should accept the above “readings” of 1, $\frac{1}{2}$ and 0. For instance, suppose we’re willing to assert $\neg DDA$ but not willing to assert $\neg DA$. Then on the reading above, we shouldn’t be willing to assert that A has semantic value less than 1. That seems a perfectly reasonable way to go; but there is an alternative way to go, also reasonably natural, according to which the commitment to $\neg DDA$ should be enough to assert that A has value less than 1. On that alternative, commitment to A not having semantic value 1 doesn’t suffice for commitment to its not being determinately true. There really isn’t a terribly deep issue of choosing between these two ways to go: it seems like basically a matter of convention for how we

understand the technical notion of semantic value. (It is a convention whose need arises only if we accept the existence of HOI_3 .)

So there are issues about the understanding of the notion of semantic value that may arise if HOI_3 is allowed for; but they don't affect the legitimacy of the semantics K_3 , and there is no need to settle them here.¹⁷ I might mention, though, that if we regard having value 1 as meaning 'determinately true' (which, I repeat, is a possible convention even given HOI_3), we should not view it as any *explanation* of the notion of indeterminacy to say: A is indeterminate if it has value other than 0 or 1. For the notion of semantic value is no clearer than the notion of indeterminacy.

4. The conditional and the determinacy operator. I've said that I'm taking K_3 to be the logic of the connectives *other than the conditional*; but K_3 is inadequate for logic generally, because it simply doesn't contain a reasonable conditional. In classical logic we use the material conditional $A \supset B$, defined as $\neg A \vee B$. It is controversial whether this representation of the conditional is adequate to our needs in a classical logic context, but it should be uncontroversial that it is hopeless as a conditional when excluded middle is gone: we don't even get the law "if A then A" on the material conditional reading, since the material conditional reading makes it equivalent to "A or not A". The lack of "if A then A" may seem not such a big deal, for how often do we bother to assert sentences of that form? But the lack is symptomatic of other problems: e.g. we also don't get the laws "if A and B, then A" or "if A, then A or B". Nor do we get "If x is tall and y is taller than x then y is tall". For these and many other reasons, it seems impossible to carry out ordinary reasoning in K_3 . What we need to do, I think, is to add a new conditional \rightarrow onto K_3 . There is some question as to exactly what the logic of the \rightarrow should be taken to be; I'll make a few remarks on this below.

But first I'll give another reason for thinking that in a nonclassical logic we need a conditional other than the material conditional \supset . Kit Fine (1975) argued in favor of using classical logic to treat vagueness, partly on the ground that nonclassical logics like K_3 couldn't handle what he called *penumbral connections* between distinct vague terms. How for instance are we to express the fact that calling an object blue is inconsistent with calling it green? The natural way to do this is to say: as a matter of conceptual necessity, anything that is blue is not green: $\Box \forall x [\text{Blue}(x) \rightarrow \neg \text{Green}(x)]$. But on the assumption that \rightarrow is \supset , we can't say this: if something is on the border between blue and green, the application of 'Blue(x) \supset \neg Green(x)' to it has value $\frac{1}{2}$. Similarly, we'd like to be able to say that something is in the region that includes blue and green and everything in between, without having to introduce a primitive notion 'blue-to-green'. (And if we did have a notion 'blue-to-green', we'd like to explain how it connects to 'blue' and to green'.) The natural way to say this is to say: if o isn't blue, it's green. But again, we can't say this on the ' \supset ' version of 'if...then': its value will be $\frac{1}{2}$, when the value of 'o is blue' and 'o is green' are $\frac{1}{2}$.

I take it that this discussion provides ample motivation for adding a new conditional \rightarrow onto the Kleene logic K_3 . But what should the logic be like? We should *not* assume that \rightarrow is truth functional in the three values 1, $\frac{1}{2}$ and 0; in fact it is easy to argue that no 3-valued truth function is remotely adequate. So the semantics of the full logic with the \rightarrow will have to be a lot more complicated than anything we've seen so far. A popular approach has been to use Lukasiewicz continuum-valued semantics, generally called “fuzzy logic”; but as remarked before, this will not do if we want to treat the semantic paradoxes as a special case of indeterminacy and maintain the intersubstitutivity of $\text{True}\langle A \rangle$ with A , for it leads to something closely akin to inconsistency. (See Restall 1992 and Hajek *et al* 2000.) I do know of one way to treat the \rightarrow that is adequate to preserving the naive theory of truth in face of the paradoxes, and it seems reasonably natural in the case of indeterminacy generally. Very roughly, we use not just a single 3-valued valuation for the atomic sentences (even for a single possible world), but a class V of them, with a certain geometric structure of relative similarity; one member v_0 of V is singled out as privileged. The value of $A \rightarrow B$ in any valuation v in V is taken to be:

- 1, if there is a valuation $v^* \neq v$ in V such that at all valuations at least as close to v as v^* is, the value of A is less than or equal to that of B ;
- 0, if there is a valuation $v^* \neq v$ in V such that at all valuations at least as close to v as v^* is, the value of A is greater than that of B ;
- $\frac{1}{2}$, if there is no valuation $v^* \neq v$ in V satisfying either of the above conditions.

The valuation of atomic predicates is subject to the following constraint: if the value of $p(o_1, \dots, o_n)$ is 1 at v_0 , then there is a valuation $v^* \neq v_0$ in V such that the value of $p(o_1, \dots, o_n)$ is 1 at all v^* that are at least as close to v_0 as v^* is; and analogously for 0, though not necessarily for $\frac{1}{2}$. This is a slight simplification of the account, leaving out a couple of complications relevant only to deal with the semantic paradoxes; for a more detailed account with simplifications removed, see Field 2003d.¹⁸ (The account can be presented as a degree-functional account in a certain space of semantic values—see Field 2003c. But that presentation is too complex to give here; let me simply say that unlike the usual semantics for “fuzzy logic”, the degrees are not linearly ordered.)

To illustrate how this works in the case of penumbral connections: a reasonable way to respect the penumbral constraints between ‘blue’ and ‘green’ would be to require that (whatever possible world is in question) there is no acceptable valuation in which there is an object o for which the value of ‘Blue(x)’ relative to o and the value of ‘Green(x)’ relative to o sum to more than 1. Since $b+g \leq 1$ implies $b \leq 1-g$, the semantics guarantees that the value of ‘ $\forall x[\text{Blue}(x) \rightarrow \neg \text{Green}(x)]$ ’ is 1 for any acceptable valuation for any world. Furthermore, if we suppose that for a blue-green object o in a world all the valuations will be such that the value of ‘Blue(x)’ relative to o and the value of ‘Green(x)’ relative to o sum to exactly 1, then ‘ $\neg \text{Blue}(x) \rightarrow \text{Green}(x)$ ’ will have value 1

for \circ , as desired.

One important feature of this semantics is that there turns out to be a rather natural way of defining a determinately operator from the \rightarrow . (I'll call this operator D rather than \mathbf{D} , to avoid pre-judging how well it corresponds to "the intuitive notion", if there is a unique such notion.) Roughly, the definition of DA is $\top \rightarrow A$, where \top is any tautology, e.g. any sentence of form $B \rightarrow B$. (Actually we need the slightly more complicated definition $A \wedge (\top \rightarrow A)$, because of the "abnormal worlds" mentioned in the previous footnote.) The result is that the value of DA in valuation v is:

- 1, if there is a valuation $v^* \neq v$ such that at all valuations at least as close to v as v^* is, the value of A is 1;
- 0, if there is a valuation $v^* \neq v$ such that at all valuations at least as close to v as v^* is, the value of A is less than 1;
- $\frac{1}{2}$, if there is no valuation v^* satisfying either of the above conditions.

And this does seem to have the right flavor for a semantics for determinacy. Clearly DA implies A and $\neg A$ implies $\neg DA$; indeed, these are "strong implications" in the sense that $DA \rightarrow A$ and $\neg A \rightarrow \neg DA$ come out as logical truths. And if "implication" just means guaranteed preservation of value 1 at the distinguished valuation v_0 rather than at all valuations, it can be shown that A implies DA (because of the constraint on the assignment of atomic predicates). But DA is still a strengthening of A , because $\neg DA$ doesn't imply $\neg A$ (from which it follows that $A \rightarrow DA$ is not a logical truth).

I'm not certain that the semantics of the conditional sketched here is ultimately the best one possible. It does seem reasonably intuitive, and it does have the virtue that it validates a logic in which, despite self-reference, we can consistently adhere to the complete intersubstitutivity of $\text{True}\langle A \rangle$ with A , even inside the scope of the conditional.¹⁹ (I know of no other remotely attractive treatment with this virtue.) In particular, if we use this semantics and define the determinately operator in it in the way suggested, then we can consistently treat not only sentences that assert that they are not true, but also those that assert that they are not determinately true, those that assert that they are not determinately determinately true, and so forth. (The 'and so forth' means 'as far through the ordinals as the iteration is expressible': since the language contains a truth predicate with which infinite conjunctions can be expressed, this is a fair way through the transfinite.)

As just hinted, the semantics allows the determinately operator to iterate non-trivially: just as DA comes out (strictly) stronger than A , so DDA is stronger than DA , $DDDA$ is still stronger, and so on.²⁰ Thus (for better or worse) this reading of 'determinately' allows for at least the possibility of higher order indeterminacy in the controversial sense (HOI_3). This possibility is realized in cases involving the semantic

paradoxes;²¹ I'm somewhat skeptical that it is realized in more ordinary cases of indeterminacy, but have no settled view. But whether or not this happens in ordinary cases, we do still have HOI_2 : we should reject excluded middle for many claims about determinateness (e.g. the first of the versions of the Falwell example in Section 3).

If we do suppose that the transfinite progression $\text{DA}, \text{DDA}, \text{DDDA}, \dots$ keeps strengthening forever (as far as it can be defined), even for ordinary indeterminacy examples, then a question arises: might we introduce a notion of "super-determinateness", stronger than each member of the transfinite progression and amounting to something like their infinite conjunction? Some may think that we indeed already have a notion of super-determinate truth, viz. 'has semantic value 1'; but that thought involves some controversial presuppositions.²² (And some may think that super-determinateness is a better candidate than \mathbf{D} for the intuitive notion \mathbf{D} of determinateness.)

I hope it is clear that if a notion of super-determinateness is intelligible, we should not assume excluded middle for the notion: for (1) we don't want to say that there is a smallest N_0 for which 'Falwell's life began by nanosecond N_0 ' is super-determinately true, and (2) the most natural way to avoid a paradox for a sentence that asserts of itself that it is not super-determinately true is to refuse to accept excluded middle for it. So if we read \mathbf{D} as "it is super-determinate that", we still have HOI_2 . How about HOI_3 ? That is, would $\neg\text{SSA}$ imply $\neg\text{SA}$, where S is a "super-determinateness operator"? I'm not sure: it's up to the person who wants a notion of super-determinacy to say.²³ If we do have such an implication—if super-determinacy ultimately collapses even though "determinacy" in the sense of \mathbf{D} doesn't—then that would seem to suggest that super-determinacy is the better candidate for the intuitive operator \mathbf{D} . If on the other hand super-determinateness keeps iterating non-trivially forever, it would seem to have no advantages over \mathbf{D} as a representation of the intuitive determinateness operator \mathbf{D} : it would generate a notion of super-super-determinateness in the same way the ordinary notion would generate a notion of super-determinateness. I will not discuss the issue further: the important point, for my purposes, a point that is independent of these speculations, is that no matter how far we go, excluded middle is never restored.

5. Where are we? A Summary and Beyond. I began the paper by noting that there is a powerful argument (essentially due to Williamson) that for any question whatever, there is a fact of the matter as to its answer. (I focused on the question of whether Jerry Falwell's life had begun by a certain moment, but I could have used any other question: e.g., the Williamson argument can be applied to "prove" that there is a fact of the matter as to whether two space-like separated points are simultaneous, apparently vindicating Lorentz over Einstein.) I noted that attempts to evade this argument while granting the law of excluded middle seem highly problematic; it looks as if to effectively block this argument, we must abandon application of the law of excluded middle to sentences we regard as "non-factual".

Unfortunately, giving up excluded middle is not enough to allow us to assert that there is no fact of the matter about certain issues: at least, not if ‘fact’ is read in a “thin” sense according to which “it is a fact that A” is equivalent to “A”; and any “thicker” sense needs explanation. For in the “thin” sense, believing that there is no fact of the matter as to whether A would amount to believing $\neg(A \vee \neg A)$, which would lead to contradiction in almost any logic; and I argued that the use of a paraconsistent logic in which the acceptance of contradictions is OK is not helpful in the present context.

But all is not lost: for we can still believe that there is no *determinate* fact of the matter. Moreover, to the extent that we do so, that is, to the extent that our belief that there is no *determinate* fact of the matter is high, our degree of belief in there being a (plain) fact of the matter is guaranteed to be low. (Reason: $P(A \vee \neg A) = P(\mathbf{GA})$, and that must be less than $1 - T$ if $P(\neg \mathbf{GA}) > T$.)²⁴

This last point (‘Moreover, ...’) is crucial, and marks a huge contrast with the use of determinateness in many classical logic approaches such as supervaluationism. The supervaluationist grants that either A or $\neg A$; so (in the “thin” sense of ‘fact’ in which ‘it is a fact that A’ says no more than ‘A’), it is either a fact that A or it is a fact that $\neg A$: there is a fact of the matter. The supervaluationist then goes on to insist that there is no *determinate* fact of the matter (or, no fact of the matter “in the thick sense”, where it is a thick fact that A iff determinately A); but it isn’t in the least evident why anyone should care. Our intuitive repugnance to the view that there is a fact of the matter as to whether the largest whole number of nanoseconds between the beginning of Ashcroft’s life and the beginning of Falwell’s is even is not lessened much by adding that that fact doesn’t have the special property of “determinateness” or “thickness” (whatever exactly that special property may be).

By contrast, on the present view determinateness is not a property unrelated to truth: having a high degree of belief in there being no determinate fact of the matter as to whether A *commits you* to having a low degree of belief in there being a (plain) fact as to whether A. In short: the supervaluationist is committed to having degree of belief 1 that there is a fact of the matter, despite his belief in indeterminacy; I, by contrast, am committed to having a very low degree of belief that there is a fact of the matter, given my belief in indeterminacy. Indeed, if I’m certain of the indeterminacy, my degree of belief that there’s a fact of the matter is 0.

Some may feel, though, that my attitude toward excluded middle is unsatisfactory. For certain instances A, I neither assert *the truth of* $A \vee \neg A$ (which would be equivalent to asserting $A \vee \neg A$) nor assert *the non-truth of* $A \vee \neg A$ (which would be equivalent to asserting $\neg(A \vee \neg A)$). “Why don’t you come clean: is it true or isn’t it?” My answer, in brief, is that if there’s no determinate fact of the matter as to whether A then there’s also no determinate fact of the matter as to whether $A \vee \neg A$, and so no determinate fact of the

matter as to whether $A \vee \neg A$ is true; and if I don't believe that there's a determinate fact as to whether $A \vee \neg A$ is true, then I *am* coming clean in refusing to assert either that it is true or that it isn't.

More fully, $A \vee \neg A$ is equivalent to $(A \vee \neg A) \vee \neg(A \vee \neg A)$ in Kleene logic: each takes value 1 when A takes value 0 or 1 and takes value $\frac{1}{2}$ when A takes value $\frac{1}{2}$. This equivalence has two consequences. The first consequence is that if I don't accept $A \vee \neg A$, I also shouldn't accept $(A \vee \neg A) \vee \neg(A \vee \neg A)$. Now for any sentence B (including the sentence $A \vee \neg A$), if I accept $B \vee \neg B$ but won't accept B or accept $\neg B$, that must be due to ignorance: I believe that there's a fact of the matter as to whether B , but don't know what it is. But if I don't accept $B \vee \neg B$, i.e. don't accept that there's a fact of the matter as to whether B , no one can legitimately complain of my refusal to accept B or accept $\neg B$; indeed, if I did accept B or accept $\neg B$, I'd be irrational in not accepting the weaker claim $B \vee \neg B$. As a special case of this, no one can legitimately complain of my refusal to accept $A \vee \neg A$ or accept $\neg(A \vee \neg A)$ when I don't accept $(A \vee \neg A) \vee \neg(A \vee \neg A)$; as I won't, when I don't accept $A \vee \neg A$. That's the first consequence. The second consequence of the equivalence of $A \vee \neg A$ to $(A \vee \neg A) \vee \neg(A \vee \neg A)$ is that if I'm in a position to accept $\neg \mathbf{D}(A \vee \neg A)$, then I'm also in a position to accept $\neg \mathbf{D}[(A \vee \neg A) \vee \neg(A \vee \neg A)]$; that is, I'm in a position to positively deny that there is a determinate fact of the matter as to whether $A \vee \neg A$, not merely to refuse to assert that there is a fact of the matter. In this case, it is all the clearer that I should not assert one way or the other on the pseudo-question of "Is $A \vee \neg A$ true or isn't it".

Although I think that in paradigm cases of indeterminacy (e.g. some of the Falwell sentences; or, I'd say, the Liar sentence) we should be unwilling to assert that the corresponding instances of excluded middle are true or that they are not true, I do think that we can unequivocally display our stance toward them: for I think that we should believe them and their negations each to degree 0. In the case of the Liar sentence I think we have an especially strong argument for believing the corresponding instance of excluded middle to degree 0: $L \vee \neg L$ leads to contradiction by an obvious argument, and we shouldn't have any positive degree of belief toward anything that we can clearly see implies a contradiction. With the Falwell sentences it is harder to give as decisive an argument; nonetheless, believing $A \vee \neg A$ for such sentences commits one to there being a fact of the matter, which seems manifestly absurd.

I conclude with some remarks on our understanding of the notion of determinacy, and also on the extent to which the notion is "objective". The first point I want to stress is that on my view, the phenomenon of uncertainty about (or even rejection of) the determinacy of an issue arises prior to having the concept of determinacy: it arises, in the first instance, in the guise of uncertainty about (or rejection of) instances of excluded middle. It is this that blocks the Williamson argument for factuality; we can block the argument even without having the concept of determinacy. The concept of determinacy is

most naturally introduced as a sort of projection of our attitudes toward excluded middle; and the central anchor of the concept is that the degree to which we believe $\mathbf{G}A$ (that is, $\mathbf{D}A \vee \mathbf{D}\neg A$) should always be the same as the degree to which we believe $A \vee \neg A$. This suffices to fix our degrees of belief only in *atomic* sentences that contain \mathbf{G} : it doesn't even fix our degrees of belief in their negations, due to higher order indeterminacy (of type 2). Nonetheless, it does give an initial fix on the concept of determinateness. We get a further fix by imposing additional laws.

Can we do better, and explicitly define \mathbf{D} in other terms? Certainly not in terms of the Kleene connectives. But the Kleene connectives are inadequate for serious reasoning: for many reasons quite independent of defining \mathbf{D} (e.g. the need to account for penumbral connections), we need a new conditional, and in the previous section I tentatively proposed a semantics for a conditional appropriate to vagueness. I also indicated a way in which we might explicitly define a determinateness operator \mathbf{D} in terms of that conditional; whether this is a fully accurate reflection of “the” intuitive notion of determinacy (assuming there is a unique such notion) is something on which I remained non-committal, but even if not, the laws governing the defined notion \mathbf{D} may be illuminating for the intuitive notion \mathbf{D} . For instance, presumably if $\mathbf{D}A$ differs from DA it is a strengthening of it, and hence $\neg(\top \rightarrow A)$ entails $\neg\mathbf{D}A$; so any grounds for denying $\top \rightarrow A$ (asserting its negation) are grounds for denying $\mathbf{D}A$. Correspondingly, any grounds for denying $(\top \rightarrow A) \vee (\top \rightarrow \neg A)$ are grounds for denying $\mathbf{G}A$. This gives a considerable handle on the problem, left open at the end of Section 2, of giving a sufficient condition on when the assertion of $\neg\mathbf{G}A$ was justified, a handle that doesn't depend at all on whether \mathbf{D} is actually definable in terms of \rightarrow .

Suppose that the intuitive notion \mathbf{D} is not definable in other terms, even including the \rightarrow . What consequences would this have for our understanding of the notion of determinacy? The consequence would be that our only understanding of this notion would come via its conceptual role. The conceptual role includes the logical laws governing \mathbf{D} , for instance the entailment from $\neg(\top \rightarrow A)$ to $\neg\mathbf{D}A$ just mentioned; it also includes laws governing degrees of belief, such as the law that $P(\mathbf{D}A \vee \mathbf{D}\neg A)$ should always be the same as $P(A \vee \neg A)$ or the essentially equivalent law that $P(\mathbf{D}A)$ should always be the same as $P(A)$. Indeed, even if \mathbf{D} is strictly definable in terms of \rightarrow , this merely shoves the question of our understanding of \mathbf{D} back to our understanding of \rightarrow ; and the understanding of that must be in terms of its conceptual role. We do have a richer set of logical laws for \rightarrow than for \mathbf{D} , so the burden on the “probability theory” is lessened. But I think that here too the “probability theory” has a role to play: for instance, since $A \wedge (\top \rightarrow A)$ is by hypothesis equivalent to $\mathbf{D}A$, we would need that the degree of belief in $A \wedge (\top \rightarrow A)$ must be the same as the degree of belief in A . So with or without a strict definition of \mathbf{D} in terms of \rightarrow , our understanding of \mathbf{D} is partly constituted by the laws governing our degrees of belief.

There may be some basis for arguing that there is a subjective component to our notion of determinacy. But there is no easy argument to this conclusion from the fact that our understanding of determinateness is partly constituted by its conceptual role, including by the laws governing our degree of belief in it. There are many notions that are objective, or arguably so, which are explained by starting from our subjective attitudes. Consider the notion of objective chance: attempts to explain this typically appeal to some form of the principle that if we attribute objective chance p to an outcome A and have no other beliefs relevant to whether A occurs then our degree of belief in A ought to be p . This subjective aspect of the concept of objective chance seems in fact to be essential to that concept. It's *possible* to regard this as showing that so-called 'objective chance' isn't really objective, but such an attitude is by no means inevitable.

It can in fact be argued that even the basic concepts of logic, such as negation and conjunction, are to be explained "subjectively", i.e. in terms of the principles governing (unconditional or conditional) degrees of belief in sentences containing them: e.g. an essential component of the concept of negation, not easily derivable from anything more fundamental, is that $P(\neg A)$ should never be more than $1 - P(A)$ (or the analogous principle involving conditional degrees of belief). Indeed, it is hard to see *how other than by invoking principles about degrees of belief* we could hope to get very far toward explaining negation, conjunction, and the like. For instance, as has often been pointed out, "explanations" like

'not A' is true if and only if A is not true

and

'A and B' is true if and only if A and B are true

are blatantly circular; they are no better at explaining these notions than

'Determinately A' is true if and only if it is determinately the case that A is true would be at explaining the notion of determinacy. But even if our understanding of notions like negation and conjunction comes entirely by the principles that govern it, including the constraints on our degrees of belief, still it is hard to believe that negation and conjunction aren't perfectly objective notions. And I see no clear reason why the case of determinateness should be any different.²⁵

Notes

1. Williamson focuses more on examples like 'Rembrandt is old' and 'Williamson is thin'; these examples have the disadvantage that 'old' and 'thin' are highly context-dependent and it isn't entirely easy to disentangle their context-dependency from their vagueness and apparent nonfactuality.

2. For an attempt to overcome this difficulty with a classical logic approach, see Field 2000 and 2003a: this attempt makes use of a nonstandard theory of propositional attitudes. For a number of inter-related reasons, I'm no longer optimistic about it.
3. Indeed, we can conclude both (i) that there are multiple nanoseconds during which his life began, rather than one, and (ii) that there is no nanosecond during which his life began.
4. The attempt to treat "non-factual discourse" within classical logic that I offered in Field 2000 also had this feature. For though it kept to classical logic, it appealed to nonclassical degrees of belief toward sentences that the agent regards as non-factual; these degrees of belief arose for the language without the determinately operator, and facilitated the later introduction of it. There is some connection between that view and the one to be offered here, as will emerge (see also the Postscript to that paper, in the reprinted version), but I don't think that the view offered in that paper really succeeded in overcoming the problem of reconciling non-factuality with excluded middle.
5. Reichenbach used the terminology of physical meaningfulness in connection with examples with a slightly different flavor from the Falwell example: e.g., he called the question of the temporal priority between two spacelike separated events physically meaningless. I think that's another compelling example of indeterminacy. Here too I have no objection to using the phrase "physically meaningless" or "metaphysically meaningless", provided one realizes it is not meaningfulness in any ordinary sense, but is just indeterminacy under another name. (And provided one does not presuppose Reichenbach's analysis of physical meaningfulness in terms of unverifiability.)
6. Besides disguising this fact, the "doesn't express a proposition" terminology also tends to obscure the possibility of certain kinds of higher order indeterminacy: for there is a strong temptation to suppose that claims of form 'Sentence S expresses a proposition' express propositions.
7. No exception needs to be made for paradoxical sentences: with excluded middle gone, the semantic paradoxes allow for the intersubstitutivity in full generality. More on this below.
8. Conversely, (8) follows from $P(A \vee \neg A)=1$, given (6) and (7). This needn't prevent a proponent of *classical* logic from adhering to non-standard probabilities that violate (8): the best way is to weaken (6). (See Field 2000, or the much earlier Shafer 1976.)

9. We could take T to be 1, but only if we are extremely generous about attributing degree of belief 1. If as I prefer we take T to be less than 1, some would argue that the lottery paradox prevents a strict identification of acceptance with degree of belief over the threshold; I doubt that it does, but to avoid having to argue the matter I have avoided any claim of strict identification.

10. This uses double negation elimination, so an intuitionist can escape *this* argument for the paradox; but there are other arguments that intuitionism fails to block.

11. So $\mathbf{G}A$ is equivalent to $\mathbf{G}\neg A$. Given this, \mathbf{D} could be defined in terms of \mathbf{G} : $\mathbf{D}A$ iff $\mathbf{G}A \wedge A$.

12. Someone willing to take the threshold of correct assertion to be 1 could make do with the weaker requirement that $P(\mathbf{D}A)=1$ iff $P(A)=1$ and $P(\mathbf{D}A)=0$ iff $P(A)=0$.

13. Moreover, this indeterminacy may “survive information about facts deemed determinate”: that is, $P_x(A \vee \neg A | C)$ may be indeterminate for every C for which $P_x(C \vee \neg C)=1$. Probably the term (HOI_1) should be restricted to this case, for it is doubtful that it intuitively counts as any sort of higher order indeterminacy otherwise. Thanks to Stephen Schiffer for pointing out the need of this qualification.

14. There are intuitions that go contrary to this (as Stephen Schiffer has emphasized to me): sometimes we seem prepared to assert A but not to assert “It is determinately the case that A ”. I’m somewhat inclined to think that this is so only in examples where A contains terms that are context-dependent as well as indeterminate, and that it is so because we give to ‘determinately A ’ a meaning like ‘under all reasonable contextual alterations of the use of these terms, A would come out true’; and this seems to me a use of ‘determinately’ different from the one primarily relevant to the theory of indeterminacy. But I confess to a lack of complete certainty on these points: for instance, another possibility would be to allow that sometimes the *upper* degree of belief plays a role in governing assertion. (My view of course straightforwardly allows that we can be willing deny the determinacy of A while being unwilling to deny A .)

15. An alternative would be to use the weaker logic K_3^- , in which the valid inferences are those in which in all valuations the value of the conclusion is at least the smallest value assumed by the premises. But this logic is not only weaker (in the sense that fewer inferences are valid since the requirements on validity are stronger), it is also less intuitive. This is especially evident when we come to add conditionals to the logic, for on reasonable treatments of the conditional $A \rightarrow B$ will sometimes get value $\frac{1}{2}$ when A gets

value $\frac{1}{2}$ and B gets value 0, and that would invalidate modus ponens if we used the stronger requirements on validity in K_3^- . In any case, the strong validity of K_3^- is definable from validity in the sense I'm taking as basic: the inference from Γ to B is strongly valid if it is valid and in addition there is a finite subset Γ_0 of Γ such that the inference from $\neg B$ to the disjunction of negations of members of Γ_0 is also valid. I don't believe that we could define validity in the sense I've taken basic in terms of strong validity, and this seems another reason to take validity rather than strong validity as the basic notion.

16. To present the generalization precisely, we must either assume a language where every object is named or else generalize from assigning values to *sentences* to assigning them to *pairs of formulas and assignment functions*.

17. Indeed, the issues aren't easily raised without assuming that the determinately operator **D** is part of the object language; and if we have the controversial sort of higher order indeterminacy, **D** can't be degree-functional on the three degrees 1, $\frac{1}{2}$ and 0, so we might want to expand the set of semantic values anyway to deal with that language.

18. For instance, to make the account adequate to Curry-like paradoxes, we need to distinguish between "normal" and "abnormal" valuations. At both, the connectives other than \rightarrow are still treated by the Kleene rules. At normal valuations v , we use the valuation rules for \rightarrow in the text. At abnormal valuations v , we use almost those rules, but in the clauses for 1 and 0 we exempt v itself from the range of valuations quantified over. (Also, it is technically convenient to require that the world v^* not be maximally far from v ; doing this allows the possibility of an abnormal world at which all conditionals get value $\frac{1}{2}$.)

19. This is proved in Field 2003b, though the semantics used there may seem on its face to be quite different from the one sketched here; see Field 2003d for the connection between the two.

20. On my weak definition of 'implies', we of course have that DA implies DDA. But we don't have that $\neg DDA$ implies $\neg DA$. And now that we have \rightarrow in the language, we can say more: we don't have that $DA \rightarrow DDA$ is a logical truth. (Incidentally, $\neg A \rightarrow D\neg DA$ isn't a logical truth either.)

21. For instance, a sentence that asserts of itself that it is not determinately true can't be said not to be determinately true, but can be said not to be *determinately* determinately true.

22. First, it presupposes that semantic value is not simply an artificial technical notion; second, it presupposes the second horn of the conventional choice discussed near the end of Section 3. It should of course be just as controversial that we understand a notion of semantic value that satisfies these presuppositions as it is that we understand a notion of super-determinateness.

23. And I repeat that the question only arises if the ordinary determinateness operator D iterates nontrivially forever.

24. Recall that because of HOI_2 , the converse fails: $P(A \vee \neg A)$ is $P(\mathbf{G}A)$, but it could be low while $P(\neg\mathbf{G}A)$ is also low.

25. This paper has gone through many versions and some change in doctrine, and a number of people made helpful comments along the way. I'd especially like to thank Joshua Schechter and Stephen Schiffer for comments at several stages in the process, and Richard Dietz for sending me a good critique of some details of the account in Field 2000 that proved highly relevant to the issues here.

References

- Field, Hartry 2000. "Indeterminacy, Degree of Belief, and Excluded Middle". *Nous* 34: 1-30. Reprinted with new Postscript in Field, *Truth and the Absence of Fact* (Oxford: Oxford University Press, pp. 278-311).
- Field, Hartry, 2003a. "Mathematical Undecidables, Metaphysical Realism, and Equivalent Descriptions". Forthcoming in Schilpp, ed., *The Philosophy of Hilary Putnam* (Open Court).
- Field, Hartry, 2003b. "A Revenge-Immune Solution to the Semantic Paradoxes". *Journal of Philosophical Logic*, forthcoming.
- Field, Hartry, 2003c. "Is the Liar Sentence Both True and False?". Forthcoming in Beall and Armour-Garb, ed., *Deflationism and Paradox* (Oxford University Press).
- Field, Hartry, 2003d. "The Semantic Paradoxes and the Paradoxes of Vagueness". Forthcoming in Beall and Glanzberg, eds., *Liars and Heaps* (Oxford University Press).
- Fine, Kit 1975. "Vagueness, Truth and Logic". *Synthese* 30: 265-300.
- Hajek, Petr, and Jeff Paris and John Sheperdson 2000. "The Liar Paradox and Fuzzy Logic". *Journal of Symbolic Logic* 65: 339-46.
- Priest, Graham 1998. "What Is So Bad About Contradictions?" *Journal of Philosophy* 95: 410-26.
- Reichenbach, Hans 1938. *Experience and Prediction*. Chicago: University of Chicago Press.
- Restall, Greg 1992. "Arithmetic and Truth in Lukaciewicz's Infinitely Valued Logic". *Logique et Analyse* 139-40: 303-12.
- Shafer, Glen 1976. *A Mathematical Theory of Evidence*. Princeton: Princeton University Press.
- Williamson, Timothy 1994. *Vagueness*. London: Routledge.