

Winner of *The Philosophical Quarterly* Essay Prize 2002

FUTURE CONTINGENTS AND RELATIVE TRUTH

BY JOHN MACFARLANE

If it is not now determined whether there will be a sea battle tomorrow, can an assertion that there will be one be true? The problem has persisted because there are compelling arguments on both sides. If there are objectively possible futures which would make the prediction true and others which would make it false, symmetry considerations seem to forbid counting it either true or false. Yet if we think about how we would assess the prediction tomorrow, when a sea battle is raging (or not), it seems we must assign the utterance a definite truth-value. I argue that both arguments must be given their due, and that this requires relativizing utterance-truth to a context of assessment. I show how this relativization can be handled in a rigorous formal semantics, and I argue that we can make coherent sense of assertion without assuming that utterances have their truth-values absolutely.

I. THE PROBLEM OF FUTURE CONTINGENTS

Suppose that the world is objectively indeterministic. In some possible futures, there is a sea battle tomorrow. In others, there is not. How should we evaluate an assertion (made now) of the sentence ‘There will be a sea battle tomorrow’?

The question is difficult to answer because we are torn between two intuitions. On the one hand, there is a strong temptation to say that the assertion is neither true nor false. After all, there are possible future histories witnessing its truth and others witnessing its falsity, with nothing to break the symmetry. I shall call this ‘the indeterminacy intuition’. On the other hand, there is a strong temptation to say that the assertion does have a definite truth-value, albeit one that must remain unknown until the future ‘unfolds’. After all, once the sea battle has happened (or not), it seems quite strange to deny that the assertion was true (or false). I shall call the thought that the assertion does have a definite truth-value ‘the determinacy intuition’.

On the face of it, these two intuitions look incompatible. No surprise, then, that standard ‘solutions’ to the problem of future contingents have

been able to save only one. Some approaches save the indeterminacy intuition, others the determinacy intuition. In §§II and III I discuss these one-sided approaches and argue that they are all unsatisfactory. A satisfactory account of future contingents must give both intuitions their due.

But how? Are they not incompatible? Only in the presence of the orthodox assumption that truth for utterances is non-relative. I shall call this assumption ‘the absoluteness of utterance-truth’. No one would deny that the truth of sentences must be relativized to a context: ‘I am cold’ has no absolute truth-value, but is true in relation to some contexts of utterance, false in relation to others. But on the orthodox view, no further relativization is called for once the context of utterance has been taken into account: the truth-value of an utterance does not depend on who is asking about it, or when. Thus if we say that an assertion of ‘There will be a sea battle tomorrow’ is neither true nor false when it is made, then we cannot allow that it might acquire a truth-value later; conversely, if tomorrow (in the midst of a sea battle) we say that the assertion has turned out to be true, we cannot say that it was neither true nor false when it was made. The truth-value of an utterance is independent of the context from which the utterance is being assessed.

Given the absoluteness of utterance-truth, then, the indeterminacy intuition and the determinacy intuition are incompatible. But an adequate account of future contingents must respect both these intuitions. In §IV I draw the obvious conclusion: we must reject the absoluteness assumption. We must relativize the truth of utterances to a *context of assessment*, and we must relativize the truth of sentences to both a context of utterance and a context of assessment. This amounts to recognizing a new kind of linguistic context-sensitivity: sentence truth can vary not just with features of the context of utterance (*u-contextuality*) but with features of the context of assessment (*a-contextuality*). It is failure to make room for this kind of context sensitivity that has left us with the traditional menu of unsatisfactory solutions to the problem of future contingents.

In the second half of the paper, I try to remedy this failure. In §V I develop a semantic framework that allows for *a-contextuality*. Within this framework, I give an account of the semantics of future contingents that respects both the indeterminacy intuition and the determinacy intuition. Then in §VI I defend *a-contextuality* against a challenge to its very coherence, first raised (with somewhat less generality) by Gareth Evans. Discharging these tasks should go some way towards convincing philosophers who have grown up with the absoluteness assumption that it is not obligatory.

II. THE INDETERMINACY INTUITION

Suppose that at some moment m_0 there is an objectively possible future history h_1 in which there is a sea battle the next day, and another h_2 in which there is no sea battle the next day (see Fig. 1).¹ These alternative histories are both objective possibilities, not just epistemic ones. It is not just that we do




Figure 1: Branching histories

not know whether or not there will be a sea battle, or even that we could not know, but that both courses of events are real possibilities. Whether the world is objectively indeterministic in this sense is, of course, a substantive scientific (and perhaps metaphysical) question. I do not here presuppose an affirmative answer to this question. All I am presupposing is that talk about the future would not be incoherent in an objectively indeterministic world. Determinism may be true, but it is not for the semanticist to say so.

Now suppose that at m_0 Jake asserts ‘There will be a sea battle tomorrow’. Is his utterance true or false? The utterance takes place at m_0 , which belongs to both h_1 and h_2 . In h_1 there is a sea battle the day after m_0 while in h_2 there is not. We may assume that nothing about Jake’s intentions picks out a particular history (h_1 or h_2). Jake may take himself to be making a claim about ‘the actual future history’, but if this means ‘the future history that includes this utterance’, then it is an improper definite description. There is no such unique history. Given that nothing about the context of utterance singles out one of the histories of which it is a part, symmetry considerations seem to rule out saying either that the utterance is true or that it is false. Thus, it seems, we must count it neither true nor false. This is the indeterminacy intuition.

There are two standard ways to capture the indeterminacy intuition in a rigorous semantics. The first, due to Łukasiewicz, is to introduce a third truth-value for future contingents (i for ‘indeterminate’) and give

¹ Here I presuppose the metaphysical picture of objective indeterminism articulated in N. Belnap *et al.*, *Facing the Future* (Oxford UP, 2001), pp. 29–32, 139–41. Moments are idealized time-slices of the universe, partially ordered by a causal–historical precedence relation ($<$) with no backward branching, and histories are maximal chains of moments. For a relativistically acceptable version of branching histories, see Belnap, ‘Branching Space-time’, *Synthese*, 92 (1992), pp. 385–434. Although it is strictly correct to say that moments are *contained in* histories, I shall sometimes talk informally of histories *passing through* moments.

three-valued truth-tables for the basic logical connectives (see Fig. 2).² But Łukasiewicz’s semantics has some implausible consequences. When ‘There will be a sea battle tomorrow’ has the value *i*, so does ‘There will be a sea battle tomorrow or there will not be a sea battle tomorrow’, even though the latter sentence is not a future contingent. It will not help to change the truth-tables or add more values: the culprit is the assumption that the connectives

\neg		\wedge	<i>t</i>	<i>i</i>	<i>f</i>	\vee	<i>t</i>	<i>i</i>	<i>f</i>	\supset	<i>t</i>	<i>i</i>	<i>f</i>
<i>t</i>	<i>f</i>	<i>t</i>	<i>t</i>	<i>i</i>	<i>f</i>	<i>t</i>	<i>t</i>	<i>t</i>	<i>t</i>	<i>t</i>	<i>t</i>	<i>i</i>	<i>f</i>
<i>i</i>	<i>i</i>	<i>i</i>	<i>i</i>	<i>i</i>	<i>f</i>	<i>i</i>	<i>t</i>	<i>i</i>	<i>i</i>	<i>i</i>	<i>t</i>	<i>t</i>	<i>i</i>
<i>f</i>	<i>t</i>	<i>f</i>	<i>f</i>	<i>f</i>	<i>f</i>	<i>f</i>	<i>t</i>	<i>i</i>	<i>f</i>	<i>f</i>	<i>t</i>	<i>t</i>	<i>t</i>

Figure 2: Łukasiewicz’s three-valued semantics

are truth-functional.³ Suppose that ‘There will be a sea battle tomorrow’ and ‘There will be an eclipse tomorrow’ both receive the same (indeterminate) truth-value. Then no matter how many truth-values there are, and no matter what

truth-tables we use for ‘not’ and ‘or’, the sentences ‘There will be a sea battle tomorrow or there will not be an eclipse tomorrow’ and ‘There will be a sea battle tomorrow or there will not be a sea battle tomorrow’ will be assigned the same truth-value. But the former is indeterminate, while the latter is true.

A more attractive approach is the supervaluational semantics due to Thomason.⁴ On this approach, an utterance is counted as true (*simpliciter*) if it is true on all possible future histories, false (*simpliciter*) if it is false on all possible future histories. Future contingents are true on some possible future histories, false on others; so they are neither true nor false. But instances of the law of excluded middle, even those whose disjuncts are future contingents, are true on all histories, and so true *simpliciter*. This approach seems to capture the indeterminacy intuition without the implausible consequences of the truth-functional approach.

III. THE DETERMINACY INTUITION

If we think about how to assign a truth-value to Jake’s utterance at m_0 , the indeterminacy intuition seems overwhelming. But now what about someone who is assessing Jake’s utterance from some point in the future? Sally is

² J. Łukasiewicz, ‘On Three-Valued Logic’, in S. McCall (ed.), *Polish Logic* (Oxford UP, 1967), pp. 16–18.

³ See A.N. Prior, ‘Three-Valued Logic and Future Contingents’, *The Philosophical Quarterly*, 3 (1953), pp. 317–26, at p. 326.

⁴ See R.H. Thomason, ‘Indeterminist Time and Truth-Value Gaps’, *Theoria*, 36 (1970), pp. 264–81.

hanging onto the mast, deafened by the roar of the cannon. She turns to Jake and says ‘Your assertion yesterday turned out to be true’. Sally’s reasoning seems unimpeachable:

Jake asserted yesterday that there would be a sea battle today
 There is a sea battle today
 So Jake’s assertion was true.

When we take this retrospective view, we are driven to assign a determinate truth-value to Jake’s utterance: this is the determinacy intuition.

But how can we give Jake’s utterance a determinate truth-value if the future is genuinely open at the time of utterance? Those who have tried to save the determinacy intuition have typically resorted to the following expedient. Out of all the possible futures at the moment of utterance, one is marked out as ‘the actual future’, as if with a ‘thin red line’ (see Fig. 3).⁵ The




Figure 3: The ‘thin red line’

thin red line is an objective feature of the context of utterance, but not an epistemically accessible one: there is no way to know which future is the marked one, except by waiting. Positing such a thin red line looks like a way to eat our cake and have it too. By supposing that there are many objectively possible future histories,

we hang on to objective indeterminism, and by positing the thin red line, we get the determinate truth-values we need for ‘retrospective’ assessments of utterances. We are not forced to say, as the supervenientist does, that assertions of future contingents are neither true nor false.

My view is that the eating precludes the having. Like Belnap and Green, I hold that positing a thin red line amounts to giving up objective indeterminism. The non-red branches in the tree are supposed to represent objectively possible futures, but their non-redness indicates precisely that they will not be the continuations of the history that includes the utterance in question. Looking down on the tree of branching histories from above, God can see that given the past and the context of utterance, only one continuation remains in play: the one marked with the thin red line. In what sense, then, are the others really ‘possibilities’? They are possible in an epistemic sense: the utterer does not know which history is marked out with the thin red line. But objectively speaking they are not genuine possibilities at all.

⁵ The metaphor is due to Nuel Belnap and Mitchell Green, ‘Indeterminism and the Thin Red Line’, in J. Tomberlin (ed.), *Philosophical Perspectives 8: Logic and Language* (Atascadero: Ridgeview, 1994), pp. 365–88, revised as ch. 6 of Belnap *et al.*, *Facing the Future*.

The idea that it makes sense to talk of a thin red line is, I think, an illusion which results from a conflation of 'external' with 'internal' perspectives in semantics. We do often say that one of two objectively possible outcomes of a past event turned out to have been the actual one, and even that it 'was going to be' the actual one. When the coin lands heads up, we can say that heads was the side that 'was actually going to land facing up', even if we do not think that the outcome was predetermined. It is this kind of talk that makes the thin red line seem intelligible and even compelling. But such talk makes sense only from some particular perspective within the tree of branching histories. From the point of view of an observer at m_1 (in Fig. 1) the actual future at m_0 held a sea battle, while from the point of view of an observer at m_2 it did not. And *qua* semanticists, we do not speak from the perspective of any particular moment on the tree of branching histories; instead we take a God's eye point of view, looking down on the tree from the outside, and try to say how the truth of sentences depends on features of the context of utterance. From this external point of view, there is no sense to saying that one of two histories passing through a moment is 'going to be the actual one'. It is only if we blur our vision, taking up internal and external perspectives simultaneously, that it can seem to make sense to mark out one of the histories in the tree (as seen from above) with a thin red line.

David Lewis sees this point very clearly. He acknowledges that if an utterance of a future contingent belongs to more than one possible future history, we cannot appeal to 'the actual future' to secure it a determinate truth-value. But Lewis does not want to give up the determinacy intuition. His solution is to reject branching altogether.⁶ On his view, each utterance takes place at a unique possible world, and each possible world has a unique future history. On these assumptions, the context of utterance always determines a unique 'actual future'.

Like the 'thin red line', however, Lewis' picture saves the determinacy intuition only by sacrificing genuine objective indeterminism. Given a context of utterance, there is only one possible future history that contains it: the future is in that sense determined. Granted, there are other possible worlds that are qualitative duplicates of the actual world up to the present and diverge thereafter, but these worlds contain different utterances (and utterers), mere 'counterparts' of the actual ones. Thus the future is open only in the sense that we do not (and perhaps cannot) know what it will bring.

⁶ See Lewis, *On the Plurality of Worlds* (Oxford: Blackwell, 1986), pp. 206–9. Lewis' arguments against branching are effectively countered by Belnap *et al.* in *Facing the Future*, pp. 206–9.

IV. THE ABSOLUTENESS OF UTTERANCE-TRUTH

To sum up, if we focus on m_0 the indeterminacy intuition seems compelling, and we are pushed towards saying that Jake's utterance is neither true nor false. But if we focus on a later moment m_1 or m_2 at which the predicted sea battle is (or is not) happening, the determinacy intuition seems compelling, and we are pushed towards saying that Jake's utterance is determinately true (or false). I have argued that traditional approaches which save only one of these intuitions at the expense of the other are inadequate. But is it possible to do better? On the face of it, there is no way to capture both. If Jake's utterance is neither true nor false, as the indeterminacy intuition demands, it is not true and it is not false. But the determinacy intuition demands that it must be one or the other.

However, this quick argument for incompatibility assumes the absoluteness of utterance-truth. If utterance-truth were relativized to the context at which the utterance is being assessed, then we could accommodate both intuitions easily. We could say that Jake's utterance is true as assessed from m_1 , false as assessed from m_2 , and neither true nor false as assessed from m_0 . In fact this is precisely what I think we should say, and what people unschooled in philosophy naturally will say. What has kept philosophers from adopting this natural solution to the problem of future contingents is their deeply entrenched theoretical commitment to the absoluteness of utterance-truth. So much the worse for absoluteness. If we need to reject it to get a plausible account of our talk about the future, then reject it we should.

I have said that philosophers of language are deeply committed to absoluteness. I have sometimes heard this questioned, on the following ground. According to one dominant paradigm in the philosophy of language, an (assertive) utterance expresses a proposition, and propositions – contingent ones, anyway – are true at some possible worlds or situations, false at others. For example, the proposition Sam expresses when he says 'I am cold' on New Year's Eve 2001, the proposition that Sam is cold on New Year's Eve 2001, is true with respect to some possible worlds (including worlds in which Sam, or a counterpart of Sam, is in Norway on New Year's Eve 2001), and false with respect to others (including worlds in which Sam, or a counterpart of Sam, is in Australia on New Year's Eve 2001). The objection goes as follows: to say that an utterance is true is to say that the proposition it expresses is true. But on the standard picture, the truth of propositions is relativized to worlds or situations. So the standard picture cannot avoid

relativizing utterance-truth to worlds as well. If this is right, then the standard picture is not committed to absoluteness after all. Even Sam's mundane utterance is true as assessed from some possible worlds, false as assessed from others. But this objection rests on a misunderstanding. Yes, the proposition expressed by Sam's utterance is true with respect to some possible worlds, false with respect to others. But to say that an utterance is true is to say more than that the proposition it expresses is true: it is to say that this proposition is true *with respect to the world at which the utterance occurs*.

A simple example will help. Let w_1 be the actual world and w_2 a world very like the actual world, except that in w_2 the dodo never became extinct. Let u_1 and u_2 be utterances, in w_1 and w_2 respectively, of the sentence 'The dodo is extinct in the year 2002'. u_1 and u_2 express the very same proposition p , and truth for p is world-relative: p is true in w_1 and false in w_2 . None the less u_1 and u_2 can be assigned absolute truth-values: u_1 is true *simpliciter* – it accurately describes the world in which it is made – while u_2 is false *simpliciter*. All that matters for the truth of the utterance u_2 is the truth-value of the proposition it expresses at w_2 ; the truth-value of this proposition at other worlds (including the actual world w_1) is simply not relevant. So relativization of propositional truth to worlds is compatible with the absoluteness of utterance-truth. The relativization of utterance-truth to a context of assessment is a different beast entirely.

V. A FRAMEWORK FOR RELATIVE TRUTH

I have argued that in order to make good sense of future contingents, we must allow the truth of utterances to be relativized to the context from which they are being assessed. The suggestion will raise some hackles: it is widely believed that there is something incoherent about relative truth. My aim in the next two sections is to put this worry to rest. In this section, I shall show how a standard framework for the semantics of indexicals can be modified to allow for relativity of truth to a context of assessment. In the next, I shall show how the modified framework can be integrated with a plausible account of assertion.

In standard indexical semantics, truth for sentences must be relativized to a context of utterance. But for technical reasons, we cannot give a direct recursive definition of 's is true at context of utterance u '. Here is a simple proof, adapted from Kaplan.⁷ The sentence 'I am here' is true at every

⁷ See D. Kaplan, 'Demonstratives', in J. Almog *et al.* (eds), *Themes from Kaplan* (Oxford UP, 1989), pp. 481–563, at pp. 508–10. See also Lewis, 'Index, Context, and Content', in S. Kanger and S. Ohman (eds), *Philosophy and Grammar* (Dordrecht: Reidel, 1980), pp. 79–100.

context of utterance.⁸ So is ‘ $2 + 2 = 4$ ’. But ‘It is always the case that I am here’ is false at (nearly) every context of utterance, whereas ‘It is always the case that $2 + 2 = 4$ ’ is true at every context of utterance. So the truth-at-a-context profile of a sentence s does not contain enough information to determine the truth-at-a-context profile of ‘It is always the case that s ’.

For technical reasons, then, the recursive clauses of a semantic theory must define not truth at a context, but truth at a *point of evaluation* (or simply ‘a point’).⁹ Points of evaluation are sequences of parameters, for example, speaker, location of utterance, time and assignment. The recursive clauses for operators can vary these parameters independently: for example, the clause for ‘it is always the case that’ shifts only the time parameter, while that for the universal quantifier shifts only the assignment parameter. So the truth-at-a-point profile of ‘ $2 + 2 = 4$ ’ differs from that of ‘I am here’: the former sentence, but not the latter, is true at a point where *speaker* = Albert Einstein, *location of utterance* = New York City, and *time* = summer solstice, 1387. (There is no context of utterance corresponding to this combination of parameters.) Unlike its truth-at-a-context profile, a sentence’s truth-at-a-point profile does determine the sentence’s contribution to the truth-at-a-point profile of complex sentences embedding it.

I call the recursive definition of *truth at a point of evaluation* the ‘semantics proper’. Of course it is the truth-at-a-context profile that we are interested in: this tells us how to evaluate assertions and other speech acts. Truth at a point is just a technical device for defining truth at a context. So we need a definition of *truth at a context of utterance* in terms of truth at a point of evaluation; I call this definition the ‘postsemantics’ (see Fig. 4). Distinguishing these two modules will help me




Figure 4:
Semantics and postsemantics

make things clearer when I add contexts of assessment.

Here is an example. Suppose the only operators in our language are tense-operators (‘it will be the case that’, ‘it was the case that’) and historical modalities (‘it is settled that’, ‘it is historically possible that’). For simplicity, suppose we have no quantifiers. Then our points will need just two parameters: moment and history. Our semantics proper will be a recursive definition of ‘ s is true at m/h ’ (where the slash indicates that m belongs to h), along the following lines:

⁸ I ignore complications raised by recorded messages and the like, which need special treatment. See S. Predelli, ‘I Am Not Here Now’, *Analysis*, 58 (1998), pp. 107–15.

⁹ The terminology is Belnap’s: *Facing the Future*, p. 142. Points of evaluation differ from Kaplan’s ‘circumstances of evaluation’ or Lewis’ ‘indices’ in that they may include both shiftable and non-shiftable parameters, e.g., both *time* and *time of utterance*.

Semantics proper for ‘it will be the case that’ and ‘it is settled that’:

‘*Will:*ϕ’ is true at m/h iff for some $m_1 > m$ on h , ϕ is true at m_1/h

‘*Sett:*ϕ’ is true at m/h iff for every h_1 through m , ϕ is true at m/h_1 .

These operators behave just as you would expect they would: ‘*Will:*’ shifts you forward along a history, while ‘*Sett:*’ quantifies over all histories passing through a given moment. But all I have given so far is a definition of truth at a point (here, a moment/history pair). A moment/history pair is not a context of utterance, so I need a further step to define truth at a context of utterance. This step is made in the postsemantics, which mediates between *truth at a point* and *truth at a context of utterance*.

Here is where the controversy begins. Which moment/history pairs are relevant to the truth of a sentence at a context of utterance? Supervaluational postsemantics looks at every point whose moment parameter is the moment of utterance:

Supervaluational postsemantics: ϕ is true [false] at a context of utterance u iff

ϕ is true [false] at every point m/h such that

m = the moment of u

h passes through m .

Thin red line postsemantics, on the other hand, supposes that context of utterance determines a unique history parameter, as well as a moment:¹⁰

Thin red line postsemantics: ϕ is true [false] at a context of utterance u iff ϕ is true [false] at every point m/h such that

m = the moment of u

h = the ‘thin red line’ at u .

I have already explained why neither of these accounts is acceptable, and I have suggested that an acceptable postsemantics will have to reject the absoluteness of utterance-truth. How, then, can we modify the framework so that utterance-truth can be relativized to the context in which the utterance is being assessed? Plainly we are going to need sentence-truth to be doubly relativized, to a context of utterance and a context of assessment. That is, we need the postsemantics to define *truth at a context of utterance and context of assessment*, instead of merely *truth at a context of utterance*. But this change in the *definiendum* of the postsemantics is the only change that is required. We can leave the semantics proper just as it is. Moreover, although the new framework will allow us to describe sentences that are *a-contextual* – sentences whose truth-values vary with the context of assessment – the framework

¹⁰ This proposal is not touched by Belnap and Green’s semantic arguments against the use of a thin red line (*Facing the Future*, pp. 160–70). It uses the very same semantics proper as Belnap and Green endorse, and appeals to the thin red line only in the postsemantics.

itself is neutral about whether there *are* any *a*-contextual sentences. If one's definition of '*s*' is true at context of utterance *u* and context of assessment *a*' makes no reference to *a*, no sentence in the language will be *a*-contextual. Thus an advocate of the absoluteness of utterance-truth has nothing to fear from the new framework itself, only from the freedom it gives us.

Taking advantage of this freedom, however, I can give a much more satisfactory postsemantics for the simple tensed language:

Double time reference postsemantics: ϕ is true [false] at a context of utterance *u* and context of assessment *a* iff ϕ is true [false] at every point *m/h* such that

m = the moment of *u*

h passes through *m* and (if the moment of *a* > *m*) through the moment of *a* as well.

The essential structural feature of this account is what Belnap calls 'double time references'.¹¹ We evaluate ϕ with respect to the moment of utterance and all of the histories passing through both it and the moment of assessment (see Fig. 5). But my use of double time references to define truth at a




Figure 5: Double time reference

context of utterance and context of assessment is different from Belnap's. Whereas I use them to define truth at a context of utterance and context of assessment, he uses them to define when an assertion counts as 'vindicated' or 'impugned'. Thus his account of assertion appeals directly to truth at a point of evaluation, and needs

to be retooled when new parameters of points are added. My approach is to interpose another layer, the postsemantics, between the semantics proper (the definition of truth at a point) and the account of assertion. The account of assertion can then appeal to a uniform notion of truth at a context of utterance and context of assessment (for details, see §VI below). Everything specific to the structure of points, and hence to the particular expressive resources of a language, is handled in the postsemantics, and 'screened off' from the theory of speech acts, which can then be developed (as it should be) in abstraction from the details of particular languages.

¹¹ See *Facing the Future*, p. 175, and 'Double Time References: Speech-Act Reports as Modalities in an Indeterministic Setting', in F. Wolter *et al.* (eds), *Advances in Modal Logic*, Vol. III (Stanford: CSLI, 2001), pp. 1–22. There is also an anticipation of double time references in Michael Dummett, *Frege: Philosophy of Language*, 2nd edn (Harvard UP, 1981), p. 395.

The double time reference postsemantics allows us to say just what we wanted to say about Jake's utterance of 'There will be a sea battle tomorrow' at m_0 (see Fig. 1). The semantics proper tells us that 'There will be a sea battle tomorrow' (s) is true at m_0/h_1 but false at m_0/h_2 . Idealizing the context of utterance (u) and context of assessment (a) as moments, the post-semantics yields the following:

At $u = m_0$ and $a = m_0$, s is neither true nor false (because we must look at both points, m_0/h_1 and m_0/h_2)

At $u = m_0$ and $a = m_1$, s is true (because we look only at m_0/h_1)

At $u = m_0$ and $a = m_2$, s is false (because we look only at m_0/h_2).

Since an utterance is true [false] with respect to a context of assessment a iff the sentence uttered is true [false] with respect to the context of utterance and a , this account implies that Jake's utterance at m_0 is true as assessed from m_1 , false as assessed from m_2 , and neither true nor false as assessed from m_0 . This is just the result I said would respect both the determinacy and indeterminacy intuitions. The price is that we must countenance a novel kind of context-sensitivity, a -contextuality. I think that this is a small price to pay for an adequate account of future contingents.

VI. MEETING EVANS' CHALLENGE

The persistence of the problem of future contingents over two millennia attests to the reluctance of philosophers to consider abandoning the absoluteness assumption. Why have we been so unwilling to give it up? Why does a -contextuality seem so outlandish?

The best diagnosis I have seen is due to Gareth Evans, who criticizes a view on which 'the evaluation of an utterance as correct or incorrect depends upon the time the *evaluation* is made'.¹² The particular view Evans criticizes is manifestly implausible, but the grounds on which he criticizes this view are general enough to apply to any view on which the truth of utterances is relativized to a context of assessment, including the view advocated here. Evans argues, in effect, that no such view is consistent with the role played by utterance-truth (or 'correctness') in our practice of making assertions. It is incoherent, he says (pp. 349–50), to suppose that a single assertion might count as 'objectively correct' at some times but not at others:

¹² 'Does Tense Logic Rest on a Mistake?', in Evans, *Collected Papers* (Oxford: Clarendon Press, 1985), pp. 343–63, at p. 348. Evans talks of the 'correctness' of utterances instead of their truth, reserving the word 'true' for the truth of sentences relative to a point of evaluation. This difference is, I think, entirely terminological.

Such a conception of assertion is not coherent. In the first place, I do not understand the use of our ordinary word 'correct' to apply to one and the same historical act at some times and not at others, according to the state of the weather. Just as we use the terms 'good' and 'bad', 'obligatory' and 'permitted' to make an assessment, once and for all, of non-linguistic actions, so we use the term 'correct' to make a once-and-for-all assessment of speech acts. Secondly, even if we strain to understand the notion 'correct-at-*t*', it is clear that a theory of meaning which states the semantic values of particular utterances solely by the use of it cannot serve as a theory of sense. If a theory of reference permits a subject to deduce merely that a particular utterance is now correct, but later will be incorrect, it cannot assist the subject in deciding what to say, nor in interpreting the remarks of others. What should he aim at, or take the others to be aiming at? *Maximum* correctness? But of course, if he knew an answer to this question, it would necessarily generate a once-and-for-all assessment of utterances, according to whether or not they meet whatever condition the answer gave.

I need to rephrase Evans' criticism using the terminology I have introduced in this paper. When we make sincere assertions, we aim to speak the truth. But if the sentence we assert is *a*-contextual, there is no non-relativized fact of the matter as to whether our assertion is true: it is true relative to some contexts of assessment, untrue relative to others. So how can we aim to speak the truth in asserting? At best we can aim to speak the truth *as assessed from such and such a context*. But the context of utterance (including our intentions in uttering the sentence) does not pick out a uniquely relevant context of assessment. If it did, then we would not need to relativize truth to a context of assessment; the context of utterance alone would provide all the information needed to get a truth-value. Perhaps we should aim at maximum truth, truth at most contexts of assessments? But in that case too we would end up with singly relativized truth, because we would be quantifying over contexts of assessment in the postsemantics. As Evans points out, any answer to the question 'What should we aim at in assertion?' will provide a way for the postsemantics to avoid serving up doubly relativized truth. The upshot seems to be that (doubly relativized) truth at a context of utterance and context of assessment is not a suitable input to an account of assertion or other speech acts. The postsemantics must tell us in what contexts of utterance a sentence is true (full stop); otherwise we can have no understanding of what someone might be aiming at in asserting it.

I think that Evans' argument neatly articulates a worry many philosophers have had about the very coherence of what I am calling *a*-contextuality, a worry that has no doubt kept the requisite semantic machinery from being developed.¹³ But at most Evans' argument shows that

¹³ Philip Percival calls Evans' argument 'the best available defence of the absolutist's claim that any relativization of token truth is incoherent': 'Absolute Truth', *Proceedings of the Aristotelian Society*, 94 (1994), pp. 189–213, at p. 198.

a-contextuality is incompatible with a particular picture of assertion, on which assertion is like a game one can either win (by speaking the truth) or lose (by speaking falsely). Why should that be our picture? When I was young, my friends and I used to play multi-player Rochambeau. In this game, whether a move counts as winning varies from opponent to opponent. A play of 'rock' will win with respect to an opponent who plays 'scissors', but lose to one who plays 'paper'. Though one cannot aim to win *simpliciter*, the game is not incoherent. It is just different from games in which winning is not relativized to opponents. Similarly, I suggest, assertions of *a*-contextual sentences, whose truth varies from one context of assessment to another, are not incoherent: they are just different from assertions of non-*a*-contextual sentences. What we need is an account that does not rule them out from the start.

Indeed, it is not obvious that 'aiming at the truth' should play any part in an account of assertion. If we aim at anything in making assertions, it is to have an effect on other people: to inform them, persuade them, amuse them, encourage them, insult them, or (often enough) mislead them. Even if we limit ourselves to sincere assertions, truth is only our indirect aim: we aim to show others what we believe, and we aim to believe what is true. If we misrepresent our beliefs but hit the truth anyway (because our beliefs are false), we have failed to make a sincere assertion, while if we miss the truth but accurately represent our beliefs, we have succeeded in making one. Perhaps belief or judgement constitutively aims at truth; assertion does not.¹⁴

What is it, then, to make an assertion? What is one doing when one asserts a sentence? One must have certain intentions and produce certain noises, but there is no assertion unless one thereby brings about a certain kind of change in normative status. One commits oneself to the truth of the sentence asserted (at its context of utterance).¹⁵ But what kind of a commitment is this? When one commits oneself to the truth of a sentence, what exactly is one committed to *doing*?

I suggest that one is committed to producing a justification, that is, giving adequate reasons for thinking that the sentence is true (relative to its context of utterance and the asserter's current context of assessment), whenever the assertion is challenged.¹⁶ (Not every objection to an assertion will count as a

¹⁴ This point of disanalogy between judgement and assertion means that my defence of relative utterance-truth may not generalize to a defence of relative judgement-truth. Cf. J. Campbell, 'The Realism of Memory', in R. Heck Jr (ed.), *Language, Thought, and Logic* (Oxford UP, 1997), pp. 157–81, at p. 165. I plan to address this intricate issue in future work.

¹⁵ Cf. J.R. Searle, 'The point or purpose of the members of the assertive class is to commit the speaker (in varying degrees) to something's being the case, to the truth of the expressed proposition': *Expression and Meaning* (Cambridge UP, 1979), p. 12.

¹⁶ The basic structure of this account of assertion comes from R. Brandom, *Making it Explicit* (Harvard UP, 1994), ch. 3, and 'Asserting', *Noûs*, 17 (1983), pp. 637–50.

challenge, in the sense at issue here. Ordinarily, the challenger must give reasonable grounds for questioning the asserter's warrant. Thus one may be justified in ignoring objections which are frivolous or unfounded, or which merely repeat challenges to which one has already responded.) If one cannot meet a challenge, whether through lack of resources or because the sentence asserted has been decisively shown to be untrue (relative to its context of utterance and one's current context of assessment), then one is obliged to withdraw the assertion. The act of withdrawal can be formal, as when a scientist retracts a claim in a journal, or informal, as when one says 'I take that back'. More often there is no explicit act of withdrawal at all, because (given natural expectations) none is needed: the speaker simply stops taking responsibility for the assertion. The norms constitutive of the practice of assertion, as I have described them, do not include an obligation to withdraw an assertion one believes or even knows to be false. Thus one can lie without violating the constitutive norms of assertion. Of course, one may be violating other, moral, norms, and if one lies too often and too egregiously, one risks no longer being treated as an asserter.¹⁷

On this account of assertion, *a*-contextuality poses no special problem. Indeed, the account allows us to describe exactly what is accomplished by the assertion of an *a*-contextual sentence. In asserting 'There will be a sea battle tomorrow' at m_0 , Jake comes to be bound by certain obligations. For example, if someone challenges the assertion at m_0 , Jake must give adequate reasons for thinking that it is true, relative to context of utterance m_0 and context of assessment m_0 . If the challenge takes the form of a conclusive demonstration that it is not yet settled whether there will be a sea battle, Jake will not be able to meet the challenge, and he will be obliged to withdraw his assertion. But if the challenge is weaker, and he meets it, then his assertion can stand. On the other hand, if at m_1 someone challenges his (original) assertion, Jake can meet the challenge by pointing to ships fighting. At this point, a proof that it was not settled at m_0 whether there would be a sea battle the next day would no longer count as a sufficient challenge to Jake's assertion, because it would not show that the sentence Jake asserted is untrue relative to the context of utterance (m_0) and his current context of assessment (m_1). Some might think it odd that a challenge that would oblige Jake to withdraw his assertion at m_0 should be ineffective at m_1 . But if there is any oddity here, it has nothing to do with *a*-contextuality. Suppose Anne makes a mathematical assertion at m_0 but does not have a proof until m_1 . Then she would be in a position at m_1 to meet a challenge she could not have met at m_0 . Same phenomenon, no *a*-contextuality.

¹⁷ Cf. Brandom, *Making it Explicit*, p. 180.

Pace Evans, then, we need not accept the absoluteness of utterance-truth in order to make sense of assertion and other speech acts. We can still think of assertions as commitments to the truth of the sentences asserted (at their contexts of utterance), for to be committed to the truth of a sentence (at a context of utterance *u*) is simply to be obliged, if challenged at any context of assessment *a*, to give adequate reasons for thinking that the sentence asserted is true (with respect to *u* and *a*), and to withdraw the assertion if the challenge cannot be met. It appears, then, that a postsemantics that defines doubly relativized truth can serve as input to a perfectly respectable account of assertion and other speech acts.

VII. CONCLUSION

I have argued that in order to give a satisfactory solution to the problem of future contingents, we need to relativize utterance-truth to a context of assessment, and sentence-truth to both a context of utterance and a context of assessment. This amounts to recognizing a new kind of linguistic context-sensitivity: in addition to being indexical in the ordinary way, or *u-contextual*, sentences can be *a-contextual*: their truth-values can vary with the context of assessment. I have shown how we can make room for *a*-contextuality in a formal semantic framework, and I have shown how this framework can be integrated with a plausible account of assertion.

Once we have accepted *a*-contextuality in sentences about the future, it is natural to look for it elsewhere. I have found fruitful applications to Lewis' theory of accommodation, epistemic contextualism, evaluative relativism, and the interpretation of our scientific predecessors' theoretical discourse. I do not think that any of these other applications *demand* *a*-contextuality, as future contingents do. In each case, there are acceptable (even if not optimal) solutions that do not require rejecting the absoluteness of utterance-truth. But once we have abandoned absoluteness and accepted *a*-contextuality in one case, there is no principled reason not to explore its applications to these other cases as well. Future contingents are important because they force us to abandon absoluteness, liberating us from its conceptual bonds elsewhere.¹⁸

University of California, Berkeley

¹⁸ This paper owes much to discussions with Nuel Belnap, Joseph Camp, Jeff King, Lionel Shapiro and Paul Teller.