

Truth in the Garden of Forking Paths*

John MacFarlane

January 28, 2007

“This network of times which approached one another, forked, broke off, or were unaware of one another for centuries, embraces *all* possibilities of time. . . . In the present one, which a favorable fate has granted me, you have arrived at my house; in another, while crossing the garden, you found me dead; in still another, I utter these same words, but I am a mistake, a ghost. . . . Time forks perpetually toward innumerable futures.”

Jorge Luis Borges, “The Garden of Forking Paths”

1 Introduction

It is very natural to think of time as a *line* running from past to future. Talk about the past concerns the part of the line that lies behind us, while talk of the future concerns the part that lies ahead. In his story “The Garden of Forking Paths,” Borges invites us to reject this linear picture and think of time as *forking* or *branching*. The branches represent possible future continuations of history. At each moment in time there are many possible continuations—many branches—and none of them is marked out as “the” future. The world itself branches.

*This paper was first presented at the LOGOS conference on Relativizing Utterance Truth in Barcelona, September 5–7, 2005. I am grateful to Richard Dietz, Iris Einheuser, Daniel López de Sa, and Max Kölbel for comments on an earlier draft, and to my Spring 2005 graduate seminar at Berkeley for helping me work through these ideas. Nuel Belnap deserves credit for getting me to think seriously about semantics in an indeterministic framework.

Unlike the linear picture, the branching picture represents the future as genuinely, objectively “open”—as contingent in a particularly strong sense. The future might be contingent in a number of weaker senses without being genuinely “open.” For example, it might be that even though there is no branching, the evolution of the world is so chaotic that it could not be predicted, even by an agent with full knowledge of its past states and all the laws of nature. Whether the future is open in a stronger sense that requires the branching picture is in part an empirical question. To see “branching worlds” taken seriously as a description of physical reality, one has only to look at the large literature on the “many worlds” interpretation of quantum mechanics.

But can we really make sense of branching worlds? On the branching picture, it seems, there is no such thing as *the* future. But we make claims about the future all the time. For example, I said ten days ago that it would be sunny today. It is sunny, so, it seems, what I said was true. But how *could* my claim have been true if, as the branching picture has it, there were both rainy and cloudy branches ahead of me when I made it?

This is the puzzle I want to discuss. Faced with the same puzzle, David Lewis concludes that we must either reject branching or accept that our future-directed talk and attitudes are fundamentally confused. (He prefers the first option.¹) I think the dilemma is a false one. In what follows, I will show that it is possible to give a semantic account of tensed talk that makes sense of our talk about the future, whether or not there is branching.

The core idea comes from my paper “Future Contingents and Relative Truth” (MacFarlane, 2003), but I hope here to correct a number of shortcomings in that paper:

1. In FCRT, I followed Prior, Thomason, Kaplan, and Belnap in treating temporal modifiers as sentential operators. This approach has now fallen out of favor among semanticists of natural language (see King (2003)). Here, I will use a more orthodox semantic framework, treating temporal modifiers as referring terms and quantifiers rather than operators.

¹See Lewis (1986, 199-209). Lewis’s main argument against branching worlds is that it “conflicts with our ordinary presupposition that we have a single future.” Lewis makes a distinction, which I do not entirely understand, between branching worlds and branching *within* a world. He concedes the possibility of the latter, though he thinks we can dismiss on “common sense” grounds the suggestion that “we ourselves are involved in branching” (209). Belnap et al. (2001, 205) rightly point out that a similar argument from common sense could be used against the claim that there are no reference-frame-independent facts about simultaneity.

2. In FCRT, I talked only of sentences and utterances, not of propositions. This made my case against supervenience easier than it should have been. In what follows, I'll show how to deal with propositions and propositional truth in a branching framework, and I'll explain how the supervenience might evade my criticism regarding retrospective assessments of future-tensed assertions by talking of proposition truth rather than utterance truth.
3. In FCRT, I did not explain how to do semantics for a one-place truth predicate in the object language. Nor did I explain how a relativist might handle an actuality operator. Both of these failings are corrected below. Indeed, it will emerge that the real problem with standard supervenience is its inability to make good sense of "actually," not its treatment of retrospective assessments of predictions.

Telling this story will require some careful concept-mongering. But with the appropriate concepts in hand, we can give a compelling semantic account of our talk of futurity and actuality that makes sense in both branching and non-branching frameworks. This should suffice to defuse arguments against branching that assume the incompatibility of branching with our ordinary talk and thought about the future.

2 A standard framework

Let's start with a basic semantic framework, as "vanilla" as possible. We can then see just what changes are necessary to accommodate branching.

Semantics in the style of Lewis (1980) and Kaplan (1989) requires relativization of truth to

- an *assignment* (to handle quantifiers),
- an *index* (to handle operators), and
- a *context* (to handle indexicals).

An *assignment* is just a function from variables to objects in the domain. (To keep things simple, we will operate with a single, fixed domain of objects.) An *index* is a collection of separately shiftable parameters. Which parameters are required depends on what operators the language contains. If the language contains modal operators, a possible world parameter is needed; if the language contains temporal

operators, a time parameter is needed; and so on. We'll work with a language that does not contain temporal operators, so for our purposes, an index can be simply a possible world. A *context* is a possible occasion on which a sentence might be used (or a representation of such an occasion). A context must be capable of supplying semantic values for indexicals and other context-sensitive expressions, but it need not be thought of as a sequence of parameters, as in Kaplan (1989).

We will call a $\langle \text{context}, \text{index}, \text{assignment} \rangle$ triple a “point of evaluation.” A semantic theory for a language consists in a recursive definition of truth at a point of evaluation for arbitrary formulas of the language. Such a theory might include clauses like the following:

- (1) $\lceil \forall \alpha \Phi \rceil$ is true at $\langle C, w, a \rangle$ iff Φ is true at every point $\langle C, w, a' \rangle$ such that a' differs from a at most in the value it assigns to α .
- (2) $\lceil \Box \Phi \rceil$ is true at $\langle C, w, a \rangle$ iff Φ is true at every point $\langle C, w', a \rangle$ such that w' is accessible from w .

When we are considering a sentence (closed formula) used in a particular context, we can strip off these relativizations and get an “absolute” truth value:

- (3) An occurrence of a sentence S in a context C is true iff S is true at $\langle C, w_C, a \rangle$, where w_C is the world of the context C and a is an arbitrary assignment. (Cf. Kaplan (1989, 522, 547).)

I will sometimes write “ S is true relative to C ” or “ S is true at C ” instead of “An occurrence of S in C is true.” All these formulations should be regarded as equivalent.²

Finally, logical truth and consequence are defined, following Kaplan (1989, 522–3), as truth and truth preservation at every context of use:³

- (4) A sentence S is logically true iff for every context C , S is true at C .
- (5) A sentence S is a logical consequence of a set Γ of sentences iff for every context C , if every member of Γ is true at C , then S is true at C .

²Note that if the language contained no indexicals, so that C had no role to play in the definition of truth at a point of evaluation, C could be omitted from the points of evaluation. It would still be needed in (3), however, to “initialize” the world index. Cf. Kaplan (1989, 594), Belnap et al. (2001, 148).

³An additional quantification over models or interpretations is required when—and *only* when—the sentences involved contain schematic “nonlogical constants.” All interpreted expressions are treated here as “logical constants.”

It is also useful to define a notion of logical necessity that quantifies over points of evaluation rather than contexts, and a corresponding notion of logical implication:⁴

- (6) A formula Φ is logically necessary iff for every point of evaluation π , Φ is true at π .
- (7) A formula Φ is logically implied by a set Γ of formulas iff for every point of evaluation π , if every member of Γ is true at π , then Φ is true at π .

(Note that these concepts, unlike logical truth and consequence, are defined for formulas in general, not just for sentences.) In many familiar semantic frameworks, a sentence is logically true iff it is logically necessary, and a sentence is a logical consequence of a set of sentences iff it is logically implied by this set of sentences. But in languages containing certain kinds of context-sensitive expressions (particularly “now” and “actually”), these notions come apart, and it is possible for a sentence to be logically true without being logically necessary, or to be a logical consequence of a set of sentences without being logically implied by them. We will see a bit later why it is important to have both sets of concepts.

3 Branching worlds

What changes do we need to make to this framework to make room for worlds that overlap in their initial temporal segments and branch towards the future? Surprisingly few. We can continue to think of indices as worlds, and the semantic clauses for temporal modifiers and indexicals like “tomorrow” can remain exactly as they were.

We’ll need some new operators to express claims about the branching structure of historical possibility from within the object language:

- (8) $\lceil Sett_t : \Phi \rceil$ is true at $\langle C, w, a \rangle$ iff Φ is true at all points $\langle C, w', a \rangle$ such that w' overlaps with w at t .
- (9) $\lceil Poss_t : \Phi \rceil$ is true at $\langle C, w, a \rangle$ iff Φ is true at some point $\langle C, w', a \rangle$ such that w' overlaps with w at t .⁵

⁴On the need for these two distinct notions, see Thomason (1970, 273) and Kaplan (1989, 548–50).

⁵Note that one can define epistemic analogues of these modalities in a framework without

These operators are called *historical* modalities because they are indexed to a time. What is settled (or inevitable) at one time may not have been settled at an earlier time, and what is still an open possibility at one time might not be later. Thus, speakers of our language can talk about the branching structure of temporal possibility by saying things like

(10) $\neg \text{Sett}_{\text{yesterday}} : \text{Sunny}(\text{today}) \wedge \neg \text{Sett}_{\text{yesterday}} : \neg \text{Sunny}(\text{today})$

It was not settled yesterday that it would be sunny today, and it was not settled yesterday that it would not be sunny today.

The only thing that *must* change in our framework when we add branching is the reference to “the world of the context” in the definition of truth for an occurrence of a sentence at a context (3). When two worlds overlap at the time of utterance, an utterance that takes place in one must take place in the other as well. So contexts, construed as possible utterance occasions, will in general be located at multiple worlds that overlap at the time of utterance. There will not, in general, be a unique “world of the context” (Belnap et al., 2001, 231–3).

Understanding this point is crucial to understanding the branching framework. It is tempting to suppose that, although the utterance takes place in many worlds or “possible histories,” they are not all on a par: one of them is marked out as “the actual history,” the one that will really come to pass. But there is no sense to be made of this idea. “Actual” and “actually,” as Lewis (1970a) argued, are indexical notions: their function is to shift the world of evaluation back to the world of the context.⁶ So we cannot appeal to the notion of actuality to single out one of the many worlds to which the context equally belongs.

What gives plausibility to the idea that one of the worlds is privileged, I think, is our tendency to picture the branching tree of histories as a branching network of roads that we are traveling down (in the back of a pickup truck, perhaps—not in the driver’s seat). We reason as follows:

genuine branching, by appealing to exact qualitative similarity of worlds up through t instead of literal overlap at t . (See Lewis (1986, 206) on “divergence.”) In this sense, to say that it is now possible that there will be a sea battle tomorrow is to say that some other possible world with a past qualitatively identical to our past (and containing “counterparts” of us) holds a sea battle tomorrow. This position amounts to a subtle form of epistemicism: it is only because we cannot now distinguish our world from worlds with qualitatively identical pasts that the future is “open” for us.

⁶This is a semantic point. It does not depend on Lewis’s modal realism. Whether we are realists or ersatzists about our worlds, we can ask whether a speaker at another world would have spoken truly in saying “Dodos are actually extinct.” It is unreasonable to suppose that the answer depends on whether there are dodos in *our* world.

Even though I'm now on *both* Route 66 and Interstate 40—they overlap here—there's a fact of the matter as to which one I'll be going down when they diverge ahead. I may not know at this point which one it will be, but I know I won't be going *both* ways! Similarly, even if I'm now located in many worlds that overlap in the present but diverge in the future, there's a fact of the matter as to which one will be my future. I'll find out when I get there.

But this picture embodies a fundamental confusion. We've already represented time as one of the spatial dimensions of our tree. So what could possibly be represented by the *motion* of a point on this tree? Certainly not a process that takes place in time, since all such processes are already represented spatially on the tree. There is nothing in the branching time model that corresponds to a car moving along the branching road, and nothing that corresponds to the decision the car will have to make to go down one branch or the other. If worlds branch, then we branch too.

4 Supervaluations

What, then, should we do with our definition of sentence truth at a context, if we can no longer talk of “the world of the context”? The obvious solution is to quantify over all of the worlds to which the possible utterance event belongs:

- (11) An occurrence of a sentence S in a context C is true iff for each of the worlds w that overlap at C and for every assignment a , S is true at $\langle C, w, a \rangle$.

This gives us a supervaluational semantics along the lines of Thomason (1970). An occurrence of

- (12) Sunny(tomorrow)
It will be sunny tomorrow

is counted true at a context just in case it is sunny on the day after the day of the context in every world to which the context belongs. Thus, if it is not settled today that it will be sunny tomorrow, (12) is not true. But it is not false either—that is, its negation is not true. So supervaluationism is committed to truth value gaps for

occurrences of sentences at contexts.⁷ This is not obviously an unwelcome result. Given that nothing in the branching picture warrants assigning one truth value as opposed to the other in such cases—the situation is completely symmetrical—one might reasonably think that future contingents *must* be “gappy.”

An attractive feature of supervaluational semantics is that it secures this result without giving up the law of excluded middle. Despite the fact that neither (12) nor its negation is logically true on the supervaluational semantics, the disjunction of (12) with its negation, that is,

(13) $\text{Sunny}(\text{tomorrow}) \vee \neg\text{Sunny}(\text{tomorrow})$

It will be sunny tomorrow or it won't be sunny tomorrow

is logically true, just as it was in the linear framework (Thomason, 1970, 270). Thus we can say things like

(14) $[\text{Poss}_{\text{now}} : \text{Sunny}(\text{tomorrow}) \wedge \text{Poss}_{\text{now}} : \neg\text{Sunny}(\text{tomorrow})]$
 $\wedge [\text{Sunny}(\text{tomorrow}) \vee \neg\text{Sunny}(\text{tomorrow})]$

Although it is now possible that it will be sunny tomorrow and possible that it will not be sunny tomorrow, either it will be sunny or it won't be.

without contradicting ourselves. If we can consistently say this in a branching universe, the branching picture does not seem as antithetical to our ordinary ways of thinking about the future as Lewis and others have supposed.

It is sometimes objected that supervaluationism secures this nice result only by making “or” non-truth-functional, which is an intolerable distortion of its meaning. I think this charge can be resisted. Truth functionality should be defined in terms of truth at a point of evaluation, not truth at a context. To see this, it is only necessary to reflect that disjunction is just as truth-functional when it appears

⁷Could the supervaluationist avoid saying this by changing the biconditional in (11) to a conditional? This alternative view would say that *S* is true at *C* if it is true at all worlds that overlap at *C*. But it wouldn't say that *S* is *not* true if it isn't true at all such worlds. In this case, the theory would just remain silent about the truth of *S* at *C*. The problem is that there are no grounds for remaining silent here. It is reasonable to withhold a verdict when the facts are still out, but that is not the case here. The theory will *still* be silent if it is if given *all* the facts about the branching worlds that could possibly be relevant to the truth of *S* at *C*. (And please don't say that there are still facts to be learned—facts about which branch will be actualized. See above on the moving car picture.) Another option (that of Belnap & Green (1994) and Belnap et al. (2001)) is to decline to assign truth values to sentences at contexts at all. Whether this option is viable depends on whether the theoretical work done by the notion of sentence truth at a context can be done in other ways. For a brief discussion, see MacFarlane (2003, 331).

between two open formulas as it is when it appears between two sentences. But open formulas do not have truth values relative to contexts, for there is no such thing, in general, as “the assignment of the context” (Belnap et al., 2001, 233–4). (That is why we *quantify* over assignments in our definition of sentence truth at a context.) So the *only* sense in which disjunction can be said to be truth-functional is this one:

- (15) A binary connective \star is *truth-functional* iff for every point of evaluation π and all formulas Φ, Ψ , the truth value of $\lceil \Phi \star \Psi \rceil$ at π is a function of the truth values of Φ and of Ψ at π .

The supervaluationist does not deny that disjunction is truth-functional in *this* sense. Perhaps there is *something* objectionable about allowing the disjunction of two sentences to be true at a context even though neither disjunct is true at that context, but it is not that doing so requires taking disjunction to be non-truth-functional. Indeed, the supervaluationist can use the very same recursive clause for disjunction as the classical logician:

- (16) $\lceil \Phi \vee \Psi \rceil$ is true at $\langle C, w, a \rangle$ iff Φ is true at $\langle C, w, a \rangle$ or Ψ is true at $\langle C, w, a \rangle$.

Philosophers have also objected that supervaluational semantics has odd *logical* upshots—for instance, that $\lceil \text{Sett}_{\text{now}} : S \rceil$ is a logical consequence of S , even though $\lceil S \supset \text{Sett}_{\text{now}} : S \rceil$ is not a logical truth. This is said to be a counterexample to the classical rule of Conditional Proof. For similar reasons, it is alleged, we must reject classical rules for Case Argument, Reductio, and Contraposition.⁸ As Williamson puts the point, “supervaluations invalidate our natural mode of deductive thinking” (Williamson, 1994, 152).

But this criticism, like the charge that supervaluationism makes disjunction non-truth-functional, seems to me misguided.⁹ Conditional Proof is a proof rule, not a semantic claim. Moreover, it is a peculiar kind of proof rule, since it makes reference to a *subproof*: it says that a conditional may be deduced from a subproof that begins with the antecedent and ends with the consequent. (Significantly, all of the rules deemed difficult for the supervaluationist are of this kind.) So the supervaluationist will only have trouble with Conditional Proof if she takes the move from S to $\lceil \text{Sett}_{\text{now}} : S \rceil$ to be legitimate in a *subproof*. But why must she do that? Whether a deductive system is sound and complete has to do with what

⁸See Machina (1976, 52-3) or Williamson (1994, 151-2), substituting “ $\text{Sett}_{\text{now}} :$ ” for “Definitely.”

⁹For kindred reflections, see McGee & McLaughlin (1998, 224–5) and 2004, 132–6.

can be *proved*, not what can be subproved. No general logical or methodological principle bars us from using a proof system that licenses certain moves only in the top level of a proof (that is, only on lines that do not depend on any hypotheses). Doing this would have its costs, of course—we could no longer simply lift proofs into subproof contexts, without checking to see if they use the special rules—but it is often the case that the addition of expressive power to a language forces us to complicate its proof theory. (Think of the extra complexities involved in moving from natural deductions in propositional logic to natural deductions in first-order logic. Indeed, even in standard first-order deduction systems, proofs cannot be indiscriminately lifted into subproofs, since flagging restrictions may be violated unless variables or individual constants are relettered.)

It may be objected that even if the supervaluationist can retain the Conditional Proof rule by restricting the inference rules that can be used in subproof contexts, such a restriction would be *ad hoc* and unmotivated. For surely

- (17) If B can be obtained from A by moves that are guaranteed to preserve truth, $\lceil A \supset B \rceil$ is guaranteed to be true.

What principled reason could there be, then, to disallow the use of a truth-preserving rule in a subproof for Conditional Proof?

But (17) is ambiguous. By “guaranteed to preserve truth,” one might mean either “guaranteed to preserve truth at a context of use” or “guaranteed to preserve truth at a point of evaluation.” So (17) might mean either

- (18) If B can be obtained from A by moves that preserve truth at every context (i.e., by logical consequences), $\lceil A \supset B \rceil$ must be true at every context (i.e., logically true).

or

- (19) If B can be obtained from A by moves that preserve truth at every point of evaluation (i.e., by logical implications), $\lceil A \supset B \rceil$ must be true at every point of evaluation (i.e., logically necessary).

The supervaluationist can perfectly well accept (19). Although the inference from S to $\lceil Sett_{\text{now}} : S \rceil$ preserves truth at every context, it does not preserve truth at every point of evaluation, and so is no counterexample to the plausible principle (17), on one natural reading.

I conclude, then, that the logical and semantic “anomalies” of supervaluationism are not so anomalous when one looks into them, and certainly no reason to reject the view. For that we must look elsewhere.

5 Retrospective assessments

In FCRT, I rejected the supervenient semantics on the grounds that it gave incorrect verdicts for *retrospective* assessments of claims about the future. Suppose that yesterday I said, “It will be sunny tomorrow,” and suppose that at the time of my utterance both sun and clouds were open possibilities for today. According to supervenientism, then, my utterance was not true. By (11), the sentence I uttered was neither true nor false at the context in which I uttered it. But surely that is the wrong verdict. I said that it would be sunny today, and look—it is sunny! How could it be, then, that what I said was not true?

To see how strange the supervenientist’s verdict is, suppose that the Director of the Bureau of Quantum Weather Prediction now offers me an irrefutable proof that, at the time of my utterance yesterday, it was still an open possibility that it would not be sunny today. Would such a proof compel me to withdraw my assertion? Hardly. If I had asserted that it was *settled* that it would be sunny today, I would have to stand corrected. But I did not assert that. I just said that it *would be* sunny—and it is. My prediction was true, as we can demonstrate simply by looking outside.

But suppose that the Director had visited me yesterday, just after I made my assertion, and confronted me with exactly the same facts. Wouldn’t I have had to acknowledge that my claim was untrue? For it would have been arbitrary to understand what I said to concern any particular one of the many worlds I occupied. (It is useless, recall, to appeal to the “actual world” in this context.) By showing that some of those overlapping worlds contained a sunny tomorrow, while others did not, the Director would have shown that there was no objective basis for calling my utterance true rather than false. I would have had to agree with the supervenientist’s verdict.

When we consider the claim just after it was made, then, the Director’s proof seems to show that it was not true. But when we consider it from a different vantage point—from the vantage point of today—the Director’s proof seems altogether irrelevant to its truth. Of course, *justification* is well known to exhibit this kind of perspectival variation. The very same considerations that count as solid grounds for belief in one context may be utterly insufficient in another, because new evidence has become available. But that’s not what we’re dealing with here. It’s not that my new evidence that it’s sunny today somehow undermines the Director’s proof that it was not settled at the time of my utterance that it would be sunny today. No, the Director’s proof still stands. I still accept its conclusion. It’s just that, from my current point of view, this proof isn’t at all relevant to the truth

of my claim.

This, then, is the puzzle:

- present claims concerning the future can be shown to be untrue by a proof of present unsettledness, but
- past claims concerning the present cannot be shown to have been untrue by a proof of past unsettledness.

6 Assessment Sensitivity

I suggested in FCRT that this odd perspective relativity was both the source of the staying power of the problem of future contingents and the key to its solution. The source of its staying power, because if one focuses on present claims concerning the future, it seems impossible that they could have a determinate truth value, while if one focuses on past claims concerning the present, it seems impossible that they could fail to have one. The key to its solution, because one has only to find a way to give both perspectives their due. The way to do that, I argued, is to relativize the truth of sentences not just to a context of use but to a context of *assessment*. One can then say that my assertion was true as assessed from the context I occupy today, neither true nor false as assessed from the context I occupied yesterday, and false as assessed from the (quite rainy) context I might have occupied today if the winds had blown a different direction.

I should emphasize that the qualifiers “of assessment” and “of use” relate to the *use* to which a context is being put semantically. They designate roles, not *kinds* of context. The very same context can be considered as context of use or as context of assessment. These roles relate to a situation in which we are assessing an utterance (actual or possible) for truth or falsity. The context of use is the context in which the utterance was made. The context of assessment is the context we occupy in assessing it.

From a technical point of view, the addition of contexts of assessment requires only minimal changes to our semantic theory. A point of evaluation can remain a ⟨context of use, index, assignment⟩ triple. Semantic clauses for expressions of the language, including the historical modalities, can also remain the same. The only real difference between the “relativist” semantics and the supervaluational semantics comes in the definition of truth for an occurrence of a sentence. We must now relativize sentence truth to *two* contexts (of use and assessment). Whereas on the supervaluational semantics, we look at all the worlds that overlap at the

context of use, on the relativist semantics, we look only at those that *also* overlap at the context of assessment—provided that context can be reached from the context of use by going forward in time along one of the branches. (If it cannot be so reached, either because it is in the past of the context of use or because it is on another branch of the tree entirely, then we just look at the worlds overlapping at the context of use.)

In stating this definition formally, it will help to have some notation:

(20) Where C is a context, $W(C)$ = the set of worlds that overlap at C .

(21) Where C_1 and C_2 are contexts,

$$W(C_1|C_2) = \begin{cases} W(C_1) & \text{if } W(C_1) \cap W(C_2) = \emptyset \\ W(C_1) \cap W(C_2) & \text{otherwise} \end{cases}$$

We can then say that

(22) An occurrence of a sentence S at C_U is true as assessed from C_A iff for every world $w \in W(C_U|C_A)$, S is true at $\langle C_U, w, a \rangle$.¹⁰

This definition gives results that accord with the intuitions elicited above. Let us consider utterances of two sentences,

(23) Sunny(tomorrow)

It will be sunny tomorrow

(24) $Sett_{\text{now}}$: Sunny(tomorrow)

It is now settled that it will be sunny tomorrow

both made by me, yesterday. Let C_1 be the context in which I uttered them and C_2 the context from which I am now assessing these utterances. Suppose that C_1 belongs to precisely four possible worlds, w_1 , w_2 , w_3 , and w_4 . Due to branching, C_2 belongs to just two of them, w_3 and w_4 . Suppose finally that in w_1 it is rainy today, while in w_2 , w_3 , and w_4 it is sunny today. (See Fig. 1.)

We then get the following (ignoring the assignment parameter of points of evaluation, since we are dealing with sentences):

$$W(C_1|C_1) = W(C_1) = \{w_1, w_2, w_3, w_4\} \quad W(C_1|C_2) = W(C_2) = \{w_3, w_4\}$$

¹⁰I will use “ S is true at (or relative to) context of use C_U and context of assessment C_A ” interchangeably with “An occurrence of S at C_U is true as assessed from C_A .”

S	Truth value ... at points of evaluation				... at contexts (use/assessment)	
	$\langle C_1, w_1 \rangle$	$\langle C_1, w_2 \rangle$	$\langle C_1, w_3 \rangle$	$\langle C_1, w_4 \rangle$	C_1/C_1	C_1/C_2
(23)	False	True	True	True	Neither	True
(24)	False	False	False	False	False	False

Notice that (23) is *assessment-sensitive*—its truth value varies with the context of assessment—while (24) is not.¹¹

Finally, we need to define logical truth and consequence. (The definitions of logical necessity and implication can remain the same as before.) Instead of quantifying over just contexts of use, we quantify over both contexts of use and contexts of assessment:

- (25) A sentence S is a logical truth iff for every pair of contexts C_1 and C_2 , S is true relative to context of use C_1 and context of assessment C_2 .
- (26) A sentence S is a logical consequence of a set Γ of sentences iff for every pair of contexts C_1 and C_2 , if every member of Γ is true relative to context of use C_1 and context of assessment C_2 , then S is true relative to context of use C_1 and context of assessment C_2 .

One consequence of this change is that $\lceil \text{Sett}_{\text{now}} : S \rceil$ is no longer a logical consequence of S . We've just seen a counterexample: at context of use C_1 and context of assessment C_2 , (23) is true but (24) is false.

7 Utterance truth and proposition truth

The argument in FCRT for taking future contingents to be assessment-sensitive rested entirely on the retrospective assessment argument I presented in the last two sections. In the next three sections, I will explain how the supervaluationist might deal with retrospective assessments. I will then consider a different sort of argument for the assessment sensitivity of future contingents.

Consider two different ways of presenting the argument from retrospective assessments:

- (27) Yesterday I uttered the sentence “It will be sunny tomorrow.”
It is sunny today.
So my utterance was true.

¹¹For this terminology, see MacFarlane (2005).

- (28) Yesterday I asserted that it would be sunny today.
It is sunny today.
So what I asserted was true.

The difference is that in (27), truth is predicated of an utterance of a sentence, while in (28), truth is predicated of “what I asserted”—a proposition.

Which of these arguments did I intend in FCRT? That is not at all clear. In the technical part of the paper, I talked of sentences and utterances, not propositions. Officially, then, the argument ought to have been something like (27). But this is not what one finds in the (less formal) part of the paper where the argument is first presented:

- (29) Jake asserted yesterday that there would be a sea battle today.
There is a sea battle today.
So Jake’s assertion was true. (325)

Here Jake’s speech act is described using indirect discourse, as in (28), by way of its propositional content rather than its sentential vehicle.¹²

I think there is a reason I slipped into proposition talk in giving the retrospective assessment argument, despite my efforts to avoid it elsewhere. I was trying to elicit the intuition that the retrospective assessment of Jake’s prediction as true was a natural one—something no ordinary person would reject. And in ordinary speech, truth and falsity are almost invariably predicated of *propositions*, as in the following:

- (30) What he said is false.
(31) Nothing George asserted in his talk is true.
(32) I know you believe he’s dishonest, but that’s false.
(33) It’s true that it has been a hot summer.
(34) That was a false claim.¹³

¹²“Jake’s assertion” in (29) denotes what Jake asserted, not Jake’s act of asserting it. Although the word “assertion” can be used to refer either to an *act* of asserting or to the content of such an act, it is doubtful that we ever predicate truth of *acts* at all, even if they are speech acts (MacFarlane, 2005, 322-3). Think about how natural it sounds to paraphrase “Jake’s assertion was true” as “What Jake asserted was true,” and how unnatural it would sound to paraphrase it as “What Jake did in asserting that was true.”

¹³Here “claim” denotes the content—what is claimed—not the act of claiming. See note 12, above.

Aside from a few, relatively isolated examples, like

(35) A truer sentence was never spoken,

people do not apply the predicates “true” and “false” to sentences or utterances, except in areas of philosophical incursion. It surprised me a bit when I realized this, because it is very common in *philosophical* prose to predicate “true” and “false” of sentences and utterances. These uses, however, must be understood as technical.

As we have seen, supervaluationism gives the “wrong” retrospective assessments of truth for past utterances of future contingents. But if I am right that utterance truth is a technical notion that plays no important role in our ordinary thought and talk, then the supervaluationist can accept these consequences without being revisionist about our ordinary future-directed talk. What really matters is whether supervaluationism can vindicate our retrospective assessments of the truth of *propositions*.

8 Propositions

To explore this question, we need to introduce ways of talking about propositional truth, both in our semantic metalanguage and in the object language.

Propositions are often said to be the “primary bearers of truth value.” What this means is that anything else that is true or false is so in virtue of expressing a proposition that is true or false. It does *not* mean that propositions have their truth values absolutely or intrinsically. In standard frameworks, propositions are taken to have truth values only relative to “circumstances of evaluation”—here, possible worlds.¹⁴ In our semantic metalanguage, then, we will have a two-place propositional truth predicate: “proposition p is true at circumstance of evaluation (or world) w .”

However, this two-place truth predicate is not what we need to evaluate ordinary retrospective assessments like “What I said yesterday was true.” “True” occurs here in the object language, as a *monadic* predicate of propositions. To make progress, we need to understand *its* semantics, which are simple and straightforward:¹⁵

¹⁴Generally, circumstances of evaluation will contain the same parameters as the indices of points of evaluation. Thus, for example, Kaplan (1989) takes them to be world/time pairs.

¹⁵Of course, once we introduce “True” into the object language, something needs to be done to ward off paradox. But this issue is orthogonal to the one we’re concerned with here, so for simplicity I’ll ignore it.

(36) “True” applies to x at a point of evaluation $\langle C, w, a \rangle$ iff x is a proposition and x is true at w .

(To avoid confusion with the various metalinguistic truth predicates in play, I’ll capitalize the monadic object-language truth predicate.)

This definition has two consequences that cohere well with our ordinary use of the truth predicate. First, it implies that every instance of the following *disquotational schema* is true at every point of evaluation (and so both logically necessary and logically true):

(37) $\forall x((x = \text{the proposition that } S) \supset (\text{True}(x) \equiv S))$ ¹⁶

Second, there is no argument place for a time in “True” as defined by (36). Thus, even if there is branching, so that worlds are eliminated as live possibilities as time passes, it is never correct to say things like

(38) What you said yesterday is True today, but it wasn’t True yesterday.¹⁷

(This is so not only on the supervaluationist’s semantics, but on the relativist’s, provided the same definition of “True” is used.) Although we do use both “is True” and “was True” as predicates of propositions, the tense does not appear to have any independent semantic significance: it is determined, rather, by the grammatical context. When a past event is being discussed, as in

(39) What you said yesterday was True,

we tend to use the past tense with the truth predicate. When a present event is being discussed, as in

(40) What you are saying is True,

we use the present tense.

Let’s use (36) to see what the supervaluationist should say about retrospective assessments of propositional truth. Yesterday I said that it would be sunny today. Today I say

(41) What I said yesterday was True.

¹⁶It is here assumed that if x is the proposition expressed by S at C , x is true at w iff S is true at $\langle C, w, a \rangle$ (for any assignment a).

¹⁷I won’t deny that we can find a use for such a sentence—but not when “what you said yesterday” denotes a proposition and literal truth is at issue.

Have I spoken truly? On the supervenient account, I have spoken truly iff (41) is true at every point $\langle C, w, a \rangle$ such that C is the current context and $w \in W(C)$. By (36), (41) is true at a point $\langle C, w, a \rangle$ iff the proposition denoted at $\langle C, w, a \rangle$ by “what I said yesterday” is true at w . But for all assignments a and all worlds $w \in W(C)$, the denotation of “what I said yesterday” at $\langle C, w, a \rangle$ is the proposition that it would be sunny today. So I have spoken truly iff the proposition that it would be sunny today is true at all the worlds that overlap at the present context. But this proposition *is* true at all such worlds (just look outside). So according to the supervenient semantics, I have spoken truly in retrospectively judging yesterday’s claim to have been true.

Here is another way to see the point. (41) is true at a point of evaluation $\langle C, w, a \rangle$ where $w \in W(C)$ iff

(42) The proposition that it would be sunny today was True

is true at that point. Since (in this framework) propositions are not tensed and propositional truth is not time-relative, (42) is true at a point iff

(43) The proposition that it is sunny today is True

is true at that point. And given the disquotational property, (43) is true at a point iff

(44) It is sunny today

is true at that point. So the supervenientist—or anyone else who employs a truth predicate defined by (36)—must assign the same truth value to (41) at C as she does to (44). Since the supervenientist takes (44) to be true at C , she must say the same about the retrospective assessment (41).

Thus, when the retrospective assessments are thought of as judgements about the truth of a proposition—as I have argued they should be—they do not pose a problem for the supervenient semantics.

9 Determinate truth

What should the supervenientist say about *contemporary* assessments of contingent claims about the future? For example, what should I have said *yesterday* about the truth of my prediction? One might expect the supervenientist to say that yesterday, after the Director’s visit, I could truly have uttered

(45) What I just said (when I said that it would be sunny tomorrow) is not True.

But the supervaluationist *can't* say this. Given the disquotational properties of "True," (45) is true at the context in question if and only if

(46) It will not be sunny tomorrow

is true at that context. Since on the supervaluationist account one cannot truly utter (46) or its negation when it is still unsettled whether it will be sunny the next day, it follows that one cannot truly utter (45) or its negation, either. Thus the semantic fact recorded in the metalanguage by the observation that neither (46) nor its negation is true at such a context is *ineffable* from the "internal" point of view. To express it, one must deploy the semanticist's technical notions of utterance truth or sentence truth relative to a context.

For those who do not think that a proof of unsettledness should compel withdrawal of an assertion about the future, this result might actually be welcome. From their "internal" points of view on the branching tree of histories, speakers will not be able to acknowledge proofs of unsettledness as grounds for asserting that what they said was not true—since these proofs would then also be grounds for asserting the negations of what they said.

But for those supervaluationists who *do* think that a proof of unsettledness should compel withdrawal of an assertion about the future, there is an easy solution. We can introduce a "determinate truth" predicate:

(47) "DetTrue" applies to x at a point of evaluation $\langle C, w, a \rangle$ iff x is a proposition and x is true at every world $w' \in W(C)$.

Using this predicate, our speakers can correctly characterize propositions whose truth is still unsettled as "not Determinately True." Whether they take a proof of unsettledness to compel withdrawal of an assertion about the future will then depend on whether they think retraction is required by a proof that the assertion is not *Determinately True*.

Like "True" and "False", "DetTrue" and "DetFalse" are not time-indexed. So it won't ever be correct to say that a proposition "is now Determinately True, but wasn't Determinately True yesterday." On the other hand, it may be that a proposition that can now be correctly called "Determinately True" could yesterday have been correctly called "not Determinately True." "DetTrue," unlike "True," is use-sensitive.¹⁸

¹⁸"True," by contrast, is assessment-sensitive (for the terminology, see MacFarlane (2005)).

It now appears that, contrary to what was alleged in FCRT, the supervaluationist *can* account for the asymmetry between contemporary and retrospective assessments of contingent claims about the future. She can acknowledge that I can now truly assert “What I said was true,” even though I couldn’t truly assert this yesterday. And she can acknowledge that I can now truly assert “What I said was determinately true,” even though yesterday I could have truly asserted “What I just said is not determinately true.” The only thing she can’t do is acknowledge today that my *utterance* yesterday was true. But the considerations raised in section 7 suggest that this need not worry us much. So the argument from retrospective assessments is not sufficient to show that we need to take the novel step of relativizing sentence truth to contexts of assessment.

Does that mean that there is no need to appeal to contexts of assessment in semantics for branching time? No. As I will now argue, the supervaluationist does not have resources for a proper treatment of “actually.” It is this that requires us to relativize truth to contexts of assessments.

10 Actuality

In standard (non-branching) frameworks, the actuality operator works as follows:

- (48) $\lceil \textit{Actually} : \Phi \rceil$ is true at $\langle C, w, a \rangle$ iff Φ is true at $\langle C, w_C, a \rangle$, where w_C is the world of the context C . (See Lewis (1970b), Kaplan (1989, 545).)

No matter how deeply embedded we are, no matter how far the world of evaluation has been shifted, the actuality operator returns it to the world of the context of use. Of course, this only works if there *is* a unique world of the context of use—as there is not when worlds can overlap and branch.

(48) respects a plausible constraint on actuality operators, which I call *Initial Redundancy*:

- (49) An operator \star is *initial-redundant* just in case for all sentences S , $\lceil \star S \rceil$ is true at exactly the same contexts of use (and assessment) as S (equivalently: each is a logical consequence of the other).

If “*Actually* :” were not initial-redundant, it might sometimes happen that you could truly utter a sentence S , but not $\lceil \textit{Actually} : S \rceil$ (or perhaps vice versa). But that does not seem to be possible. When you can truly say, “It will be sunny tomorrow,” you can truly say, “It will actually be sunny tomorrow,” and when you can truly say, “It will actually be sunny tomorrow,” you can truly say, “It will be

sunny tomorrow.” This is not because “actually” has no effect on truth conditions, but because of a delicate relation between the semantics for “actually” and the definition of sentence truth at a context. The effect of adding an actuality operator to the *front* of a sentence is to shift the world of evaluation to the world of C . This has an effect on the sentence’s truth-at-points profile, but not on its truth-at-contexts profile, because (in standard, non-branching frameworks) a sentence is true at a context C just in case it is true at the point $\langle C, w_C \rangle$, where w_C = the world of C . (Since we’re dealing with sentences, I omit the assignment parameter here and in the following paragraph.)

In a branching framework, we no longer have a unique “world of the context of use,” so we need to generalize (48).¹⁹ Here we are tightly constrained by Initial Redundancy, which requires that for all sentences S and contexts C ,

(50) S is true at C iff $\ulcorner \text{Actually} : S \urcorner$ is true at C .

To respect this constraint, the supervaluationist, who takes S to be true at C just in case it is true at $\langle C, w' \rangle$ for all $w' \in W(C)$, must say that

(51) $\ulcorner \text{Actually} : \Phi \urcorner$ is true at $\langle C, w \rangle$ iff Φ is true at $\langle C, w' \rangle$ for all $w' \in W(C)$.²⁰

On the other hand, the relativist, who takes S to be true at context of use C_U and context of assessment C_A just in case it is true at $\langle C_U, w' \rangle$ for all $w' \in W(C_U|C_A)$, must define the actuality operator as follows:

(52) $\ulcorner \text{Actually} : \Phi \urcorner$ is true at $\langle C_U, C_A, w \rangle$ iff Φ is true at $\langle C_U, C_A, w' \rangle$ for all $w' \in W(C_U|C_A)$.²¹

¹⁹The first attempt to do this of which I am aware is Belnap et al. (2001, 246–7). Belnap defines two actuality operators. The first, “*Actually*₁,” is essentially (51), which is initial-redundant in supervaluational semantics but not in relativist semantics. The second, “*Actually*₂,” which is initial-redundant for the relativist as well, will be discussed briefly below (footnote 20). Belnap is neither a supervaluationist nor a relativist; he declines to give truth values to sentences at contexts at all.

²⁰Initial Redundancy would also be secured by the clause: $\ulcorner \text{Actually} : \Phi \urcorner$ is true at $\langle C, w \rangle$ iff Φ is true at $\langle C, w \rangle$. But this makes the actuality operator redundant in all contexts, not just initially. It may be that there is a use of “actually” in English that behaves this way—Lewis’s (1970b) “non-shifty” use—but we’re after an operator that makes a difference in embedded contexts. Belnap’s “*Actually*₂” operator is a kind of hybrid between (51) and the redundant operator: it behaves like one or the other, depending on whether it is evaluated on a history that contains the moment of the context of use.

²¹Note that the addition of an actuality operator has forced us to bring the context of assessment into our points of evaluation, where it was not needed before.

So the relativist's different definition of sentence truth at a context (or contexts) requires a substantially different definition of truth at a point of evaluation for the actuality operator.²²

To see why this matters, let us reflect on the upshots of the previous three sections. From an “external” point of view—the point of view of the semanticist assigning truth values to sentences relative to contexts of use and contexts of assessment—there is a clear difference between the supervaluational semantics of (11) and the relativist semantics of (22). But from an “internal” point of view—the perspective of speakers within the tree of worlds, making claims and commenting on the truth of other claims—this difference is effectively invisible. Provided that “actually” is not in the language, the supervaluationist and the relativist can use exactly the same recursive definition of truth at a point of evaluation (even for the predicate “True” itself). The differences between their positions arise only when we need to define truth at a context of use (and context of assessment). But these are rarified metatheoretic notions: the folk assess claims for truth and falsity from *within* the language, by saying things like “That’s false” or “What she said was true.” If we stick resolutely to this internal point of view, the supervaluational theory looks indistinguishable from the relativist theory.

In FCRT, I sought to connect the theoretical notion of truth relative to a context of use and context of assessment to the participants’ linguistic practice by giving it a role in a normative account of assertion. To assert a sentence, I said, is *inter alia* to commit oneself to responding to appropriate challenges, by providing grounds for the truth of the sentence relative to the original context of utterance and the context of assessment one occupies at the time of the challenge. This account gives “internal” significance to the difference between supervaluationism and relativism; if it is correct, then the two views would have different implications about precisely what one is committing oneself to in asserting that it will be sunny tomorrow. But is it correct? One might argue that if there are norms for assertion, they should be formulable from the participants’ internal perspective, using a monadic propositional truth predicate, not the semanticist’s doubly relativized truth predicate. For example: withdraw an assertion if what you asserted is

²²The introduction of an actuality operator brings with it “anomalies” similar to those discussed at the end of section 4: for example, $\lceil \text{Actually} : S \rceil$ is a logical consequence of S , though $\lceil S \supset \text{Actually} : S \rceil$ is not a logical truth. But are these anomalies really unwelcome? After all, we designed our actuality operator precisely to ensure that $\lceil \text{Actually} : S \rceil$ was a logical consequence of S and vice versa (Initial Redundancy). If we want our language to contain an actuality operator, then, we are going to have to get over our qualms about these logical “anomalies.” Fortunately, these qualms are misplaced (see the end of section 4).

shown to have been Untrue (or, alternatively, Determinately Untrue). And, as we have seen, these norms won't distinguish between supervenational and relativist semantics.

By adding an actuality operator to the language, however, we have dragged the "external" differences between the two views down to the "internal" level. In order to secure Initial Redundancy, the supervenationalist and the relativist must give different definitions of truth at a point of evaluation for the actuality operator. These differences will manifest themselves in different verdicts about the correctness of judgements of the form, "What she said was True."

Take, for example, my claim yesterday that it would be sunny today—henceforth, my "first claim." As we have seen, the supervenationalist can agree with the relativist that it is correct, today, to say "What I said yesterday was True," even though it wouldn't have been correct, yesterday, to say "What I just said is True." But suppose that yesterday I also said

(53) *Actually* : Sunny(tomorrow)

It will actually be sunny tomorrow.

Call this my "second claim." According to the supervenationalist, it should be correct for me to say (now) that my first claim was True and my second claim False. On the relativist's account, on the other hand, I can correctly say that both claims were True. For on the relativist's account, "actually" tells us to look at the worlds that overlap at the context of *assessment*, not the context of use, and at all such worlds, it is sunny today. The relativist's verdict here more closely matches ordinary usage.

Suppose that I also made a third claim yesterday:

(54) *Actually* : Sunny(tomorrow) \vee *Actually* : Cloudy(tomorrow).

Either it will actually be sunny tomorrow or it will actually be cloudy tomorrow.

The supervenationalist and the relativist agree that *yesterday* I could not have correctly described this claim as "True." However, they differ about the *retrospective* assessment. According to the relativist, it is correct today to say that this third claim was "True." The supervenationalist must deny this. Here again, I think the relativist's view accords better with common sense.

Finally, consider an assertion, made yesterday, of

(55) $Poss_{today} : \exists x \exists y (x \neq y \wedge \text{Weather}(\text{tomorrow}) = x \wedge \text{Actually} : \text{Weather}(\text{tomorrow}) = y)$

Today it is still possible that the weather tomorrow will be different than it actually will be.

On the supervenient semantics, the embedded clause

(56) $\text{Actually} : \text{Weather}(\text{tomorrow}) = y$

will be false at every point of evaluation that includes the context of use (assuming that the weather is, at that point, still unsettled). So the whole sentence will be false at every point of evaluation that includes that context. Thus the proposition expressed by the sentence at that context can never be correctly said to be “True”—not even later, from the vantage point of a sunny tomorrow. On the relativist semantics, by contrast, an occurrence of (55) yesterday is true as assessed retrospectively from a sunny today, and from that vantage point speakers of the language can correctly characterize the proposition it expressed as “True.” Here again, the relativist’s predictions accord with common sense, while the supervenientist’s do not.

It is often supposed that a resolute commitment to a branching worlds picture would require us to give up talk of actuality. That is the dilemma Lewis offers us. He picks common sense; others pick branching (say, on the grounds that it provides the best interpretation of current physical theory). The supervenient semantics for branching does saddle us with Lewis’s dilemma, but it is not the only option. As I have shown, it is possible to give a semantics for “actually” that makes good sense of our thought and talk about the future and past, even if we live in a “garden of forking paths.” But to do so, we need to relativize truth to contexts of assessment.

References

- N. Belnap & M. Green (1994). ‘Indeterminism and the Thin Red Line’. *Philosophical Perspectives* 8:365–388.
- N. Belnap, et al. (2001). *Facing the Future: Agents and Choices in Our Indeterministic World*. Oxford University Press, Oxford.
- J. L. Borges (1964). *Labyrinths: Selected Stories and Other Writings*. New Directions, New York.

- D. Kaplan (1989). 'Demonstratives: An Essay on the Semantics, Logic, Metaphysics, and Epistemology of Demonstratives and Other Indexicals'. In J. Almog, J. Perry, & H. Wettstein (eds.), *Themes from Kaplan*, pp. 481–566. Oxford University Press, Oxford.
- J. C. King (2003). 'Tense, Modality, and Semantic Values'. *Philosophical Perspectives* **17**:195–245.
- D. Lewis (1970a). 'Anselm and Actuality'. *Noûs* **4**:175–188.
- D. Lewis (1970b). 'General Semantics'. *Synthese* **22**:18–67. Reprinted in Lewis 1983.
- D. Lewis (1980). 'Index, context, and content'. In S. Kanger & S. Öhman (eds.), *Philosophy and Grammar*. Reidel, Dordrecht.
- D. Lewis (1986). *On the Plurality of Worlds*. Blackwell, Oxford.
- J. MacFarlane (2003). 'Future Contingents and Relative Truth'. *Philosophical Quarterly* **53**:321–336.
- J. MacFarlane (2005). 'Making Sense of Relative Truth'. *Proceedings of the Aristotelian Society* **105**:321–339.
- K. F. Machina (1976). 'Truth, belief and vagueness'. *Journal of Philosophical Logic* **5**:47–78.
- V. McGee & B. McLaughlin (1998). 'Review of Timothy Williamson, *Vagueness*'. *Linguistics and Philosophy* **21**:221–235.
- V. McGee & B. P. McLaughlin (2004). 'Logical Commitment and Semantic Indeterminacy: A Reply to Williamson'. *Linguistics and Philosophy* **27**:221–235.
- R. H. Thomason (1970). 'Indeterminist Time and Truth-Value Gaps'. *Theoria* **36**:264–281.
- T. Williamson (1994). *Vagueness*. Routledge, London and New York.

Figure 1:

